Fabric OS

Administrator's Guide

Supporting Fabric OS v5.2.0

Publication Number: 53-1000239-01
Publication Date: 09/29/2006

Copyright © 2006, Brocade Communications Systems, Incorporated.

ALL RIGHTS RESERVED.

Brocade, the Brocade B weave logo, Fabric OS, File Lifecycle Manager, MyView, Secure Fabric OS, SilkWorm, and StorageX are registered trademarks and Tapestry is a trademark of Brocade Communications Systems, Inc., in the United States and/or in other countries. All other brands, products, or service names are or may be trademarks or service marks of, and are used to identify, products or services of their respective owners.

FICON, IBM @server BladeCenter are registered trademarks of IBM Corporation in the U.S. and other countries.

Notice: The information in this document is provided “AS IS,” without warranty of any kind, including, without limitation, any implied warranty of merchantability, noninfringement or fitness for a particular purpose. Disclosure of information in this material in no way grants a recipient any rights under Brocade's patents, copyrights, trade secrets or other intellectual property rights. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use.

The authors and Brocade Communications Systems, Inc. shall have no liability or responsibility to any person or entity with respect to any loss, cost, liability, or damages arising from the information contained in this book or the computer programs that accompany it.

Notice: The product described by this document may contain “open source” software covered by the GNU General Public License or other open source license agreements. To find-out which open source software is included in Brocade products, view the licensing terms applicable to the open source software, and obtain a copy of the programming source code, please visit *http://www.brocade.com/support/oscd*.

Export of technical data contained in this document may require an export license from the United States Government.

Brocade Communications Systems, Incorporated

Corporate Headquarters
Brocade Communications Systems, Inc.
1745 Technology Drive
San Jose, CA 95110
Tel: 1-408-333-8000
Fax: 1-408-333-8101
E-mail: info@brocade.com

Asia-Pacific Headquarters
Brocade Communications Singapore Pte. Ltd.
9 Raffles Place
#59-02 Republic Plaza 1
Singapore 048619
Tel: +65-6538-4700
Fax: +65-6538-0302
Email: apac-info@brocade.com

European and Latin American Headquarters
Brocade Communications Switzerland Sàrl
Centre Swissair
Tour A - 2ème étage
29, Route de l'Aéroport
Case Postale 105
CH-1215 Genève 15
Switzerland
Tel: +41 22 799 56 40
Fax: +41 22 799 56 41
Email: emea-info@brocade.com

Document History

The following table lists all versions of the *Fabric OS Administrator's Guide*.

Document Title	Publication Number	Summary of Changes	Publication Date
Fabric OS Procedures Guide	53-0000518-02	First released edition.	April 2003
Fabric OS Procedures Guide	53-0000518-03	Revised for Fabric OS v4.2.0.	December 2003
Fabric OS Procedures Guide	53-0000518-04	Revised to include switch-specific information.	March 2004
Fabric OS Procedures Guide	53-0000518-05	Revised for Fabric OS v4.4.0.	September 2004
Fabric OS Procedures Guide	53-0000518-06	Revised to add RADIUS and SSL procedures.	October 2004
Fabric OS Administrator's Guide	53-0000518-07	Revised book title. Added information about 200E, 4012, and 48000 switches.	April 2005
Fabric OS Administrator's Guide	53-1000043-01	Revised for Fabric OS v5.1.0.	January 2006
Fabric OS Administrator's Guide	53-1000043-02	Removed SilkWorm 4016 and 4020 from supported switches; FCIP chapter updates.	June 2006
Fabric OS Administrator's Guide	53-1000239-01	Revised for Fabric OS v5.2.0 features. Added new hardware platforms: SilkWorm FC4-48 and FC4-16IP.	September 2006

Contents

About This Document

How This Document Is Organized xxiii
Supported Hardware and Software xxiv
What's New in This Document xxv
Document Conventions xxvi
Additional Information xxvii
Getting Technical Help xxx
Document Feedback xxxi

Section I Standard Features

Chapter 1 Introducing Fabric OS CLI Procedures

About Procedural Differences 1-1
Scope and References 1-2
About the CLI 1-2
Help Information 1-3
Displaying Command Help 1-3
Displaying Additional Help Topics 1-4

Chapter 2 Performing Basic Configuration Tasks

Connecting to the Command Line Interface 2-2
Using Telnet or SSH session 2-2
Using a console session on the serial port 2-3
Setting the Default Account Passwords 2-4
Changing Default Account Passwords at Login 2-4

Configuring the Ethernet Interface 2-6
Displaying the Network Interface Settings 2-6
Setting Static Ethernet Addresses 2-7
Configuring DHCP 2-7
Setting the Date and Time 2-9
Setting Time Zones 2-10
Synchronizing Local Time 2-13
Maintaining Licensed Features 2-14
Customizing Switch Names 2-16
Customizing the Chassis Name 2-17
Working With Domain IDs 2-17
Activating Ports on Demand 2-19
Disabling and Enabling Switches 2-20
Disabling and Enabling Ports 2-21
Making Basic Connections 2-22
Connecting to Devices 2-22
Connecting to Other Switches 2-22
Linking Through a Gateway 2-23
Checking Status 2-24
Tracking and Controlling Switch Changes 2-25
Configuring the Audit Log 2-29
Auditable Event Classes 2-30
Shutting Down Switches and Directors 2-32
High Availability of Daemon Processes 2-33

Chapter 3 Managing User Accounts

Overview 3-1
Accessing the Management Channel 3-2
Using Role-Based Access Control (RBAC) 3-2
Configuring the Authentication Model 3-6

Managing the Local Database User Accounts 3-7
About the Default Accounts 3-7
Defining Local User Accounts 3-7
Recovering accounts 3-10
Changing Local Account Passwords 3-11
Configuring the Local User Database 3-12
Distributing the Local User Database 3-12
Protecting the Local User Database from Distributions 3-13
Configuring Password Policies 3-13
Managing Fabric OS Users on the RADIUS Server 3-17
Creating Fabric OS User Accounts 3-17
Setting Up RADIUS AAA Service 3-21
Configuring the RADIUS Server 3-22
Configuring RADIUS Servers on the Switch 3-27
Enabling and Disabling Local Authentication as Backup 3-30
Setting the Boot PROM Password 3-30
With a Recovery String 3-30
Without a Recovery String 3-33
Recovering Forgotten Passwords 3-35

Chapter 4 Configuring Standard Security Features

Secure Protocols 4-1
Ensuring Network Security 4-3
Configuring the Telnet Interface 4-4
Blocking Listeners 4-5
Accessing Switches and Fabrics 4-5
Port Configuration 4-6

Configuring for the SSL Protocol 4-7
Browser and Java Support 4-7
Summary of SSL Procedures 4-7
Choosing a Certificate Authority 4-8
Generating a Public/Private Key 4-8
Generating and Storing a CSR 4-9
Obtaining Certificates 4-10
Installing a Switch Certificate 4-10
Activating a Switch Certificate 4-11
Configuring the Browser 4-11
Installing a Root Certificate to the Java Plug-in 4-12
Displaying and Deleting Certificates 4-13
Troubleshooting Certificates 4-13
Configuring for SNMP 4-14
Setting the Security Level 4-15
Using the **snmpConfig** Command 4-16
Using Legacy Commands for SNMPv1 4-19
Configuring Secure File Copy 4-23

Chapter 5 Configuring Advanced Security

About Access Control List (ACL) Policies 5-2
How the ACL Policies are Stored 5-2
Identifying Policy Members 5-2

Configuring ACL Policies 5-3
Displaying ACL Policies 5-4
Configuring a DCC Policy 5-4
Creating an SCC Policy 5-7
Saving Changes to ACL Policies 5-8
Activating Changes to ACL Policies 5-8
Adding a Member to an Existing Policy. 5-9
Removing a Member from a Policy 5-9
Deleting a Policy. 5-10
Aborting All Uncommitted Changes 5-10
Distributing the Policy Database 5-11
Configuring the Database Distribution Settings 5-12
Distributing ACL Policies to other Switches 5-13
Setting the Consistency Policy Fabric-Wide 5-14
Notes on Joining a Switch to the Fabric 5-16

Chapter 6 Maintaining Configurations

Displaying Configuration Settings 6-1
Backing Up a Configuration 6-2
Troubleshooting Configuration Upload 6-4
Restoring Switch Information. 6-4
Restoring a Configuration. 6-5
Configuration Download Without Disabling a Switch 6-5
Security Considerations 6-7
Troubleshooting Configuration Download 6-8
Restoring Configurations in a FICON Environment 6-9
Downloading Configurations Across a Fabric 6-9
SilkWorm 48000 Configuration Form 6-10

Chapter 7 Managing Administrative Domains

About Administrative Domains 7-1
Admin Domain Features 7-3
Requirements for Admin Domains 7-4
User-Defined Administrative Domains 7-4
System-Defined Administrative Domains 7-4
Admin Domain Access Levels 7-6
Admin Domains and Login 7-7
Admin Domain Member Types 7-8
Admin Domains and Switch WWN 7-9
Admin Domain Compatibility and Availability 7-11
Managing Admin Domains 7-12
Understanding the AD Transaction Model 7-12
Implementing Admin Domains 7-13
Creating an Admin Domain 7-13
Assigning a User to an Admin Domain 7-15
Activating and Deactivating Admin Domains 7-16
Adding and Removing Admin Domain Members 7-17
Renaming an Admin Domain 7-18
Deleting an Admin Domain 7-19
Deleting All User-Defined Admin Domains 7-19
Validating an Admin Domain Member List 7-20
Using Admin Domains 7-21
Using CLI Commands in an AD Context 7-21
Executing a Command in a Different AD Context 7-22
Displaying an Admin Domain Configuration 7-22
Switching to a Different Admin Domain Context 7-23
Performing Zone Validation 7-24
Admin Domain Interactions 7-24

Chapter 8 Installing and Maintaining Firmware

About the Firmware Download Process 8-2
Upgrading and Downgrading Firmware 8-2
Effects of Firmware Changes on Accounts and Passwords 8-3
Considerations for FICON CUP Environments 8-3
Preparing for Firmware Downloads 8-3
Checking Connected Switches 8-5
Obtaining and Decompressing Firmware 8-6
Performing Firmware Download on Switches 8-7
Overview of the Firmware Download Process on Switches. 8-7
Performing Firmware Download on Directors 8-9
Overview of the Firmware Download Process on Directors 8-9
Testing and Restoring Firmware on Switches 8-14
Testing and Restoring Firmware on Directors 8-15
Validating Firmware Download 8-19
Troubleshooting Firmware Download 8-20
Downgrading Firmware from Fabric OS v5.2.0. 8-20
Blade Troubleshooting Tips 8-30
Synchronizing Firmware Versions on Partitions 8-30

Chapter 9 Configuring SilkWorm Directors

Identifying Ports 9-1
By Slot and Port Number 9-2
By Port Area ID 9-2
By Index 9-3

Basic Blade Management 9-7
Powering Port Blades Off and On 9-7
Disabling and Enabling Port Blades 9-8
Conserving Power 9-9
Blade Terminology and Compatibility 9-10
CP Blades 9-11
Port Blade Compatibility 9-11
Setting Chassis Configuration Options 9-12
Obtaining Slot Information 9-13
Configuring a New SilkWorm 24000 with Two Domains 9-14
Converting an Installed SilkWorm 24000 to Support Two Domains 9-15
Setting the Blade Beacon Mode 9-17

Chapter 10 Routing Traffic

About Data Routing and Routing Policies 10-1
Specifying the Routing Policy 10-2
Assigning a Static Route 10-3
Specifying Frame Order Delivery 10-3
Using Dynamic Load Sharing 10-4
Viewing Routing Path Information 10-5
Viewing Routing Information Along a Path 10-7

Chapter 11 Using the FC-FC Routing Service

Supported Platforms 11-1
Fibre Channel Routing Concepts 11-2
Front Domain Consolidation 11-5
Supported Configurations and Platforms 11-6
Using Front Domain Consolidation 11-6
Proxy Devices 11-8
Routing Types 11-9
Fibre Channel NAT and Phantom Domains 11-10

Setting Up the FC-FC Routing Service 11-11
Performing Verification Checks 11-11
Assigning Backbone Fabric IDs 11-13
Configuring FCIP Tunnels (Optional) 11-14
Configuring FC-FC Routing to Work with Secure Fabric OS (Optional) 11-14
Configuring an Interfabric Link 11-16
Configuring LSANs and Zoning. 11-22
Optional Configuration Procedures 11-26
FC Router Port Cost (Optional) 11-26
Setting a Proxy PID 11-29
Matching Fabric Parameters 11-29
EX_Port Frame Trunking (Optional) 11-30
Supported Configurations and Platforms 11-30
Using EX_Port Frame Trunking 11-31
Monitoring Resources 11-33
Routing ECHO 11-36
Upgrade and Downgrade Considerations 11-37
Interoperability with Legacy FCR Switches. 11-37
Backward Compatibility 11-37
Connecting to McDATA SANs. 11-38
Supported Platforms and Configurations 11-38
Configuring the Fabrics for Interconnectivity 11-39
Connectivity Modes 11-39
Configuring the FC Router 11-40
Configuring McDATA for Interconnection 11-44
LSAN Zoning with McDATA. 11-46
Completing the Configuration 11-48

Chapter 12 Administering FICON Fabrics

Overview of Fabric OS Support for FICON 12-1
Supported Switches 12-3
Types of FICON Configurations 12-3
Control Unit Port (CUP) 12-4
FICON Commands 12-4
Security Considerations 12-5
Configuring Switches 12-6
Preparing a Switch 12-7
Configuring a Single Switch 12-7
Configuring a High-Integrity Fabric 12-7
Setting a Unique Domain ID 12-9
Displaying Information 12-10
Link Incidents 12-10
Registered Listeners 12-10
Node Identification Data 12-10
FRU Failures 12-10
Swapping Ports 12-11
Clearing the FICON Management Database 12-11
Using FICON CUP 12-12
Setup Summary 12-12
Enabling and Disabling FICON Management Server Mode 12-13
Displaying the fmsmode Setting 12-14
Displaying Mode Register Bit Settings 12-15
Setting Mode Register Bits 12-16
Persistently Enabling/Disabling Ports 12-16
Port and Switch Naming Standards 12-17
Adding and Removing FICON CUP Licenses 12-17
Zoning and PDCM Considerations 12-18
Zoning and Link Incident Reporting 12-18
Backing Up and Restoring Configurations 12-18

Troubleshooting 12-19
Identifying Ports 12-20
Backing Up FICON Files 12-20
Recording Configuration Information 12-22

Chapter 13 Configuring the Distributed Management Server

Enabling and Disabling the Platform Services 13-1
Controlling Access 13-2
Configuring the Server Database 13-6
Controlling Topology Discovery 13-7

Chapter 14 Working With Diagnostic Features

About Fabric OS Diagnostics 14-1
Viewing Power-On Self Test. 14-2
Viewing Switch Status 14-4
Viewing Port Information 14-6
Viewing Equipment Status 14-9
Viewing the System Message Log 14-11
Viewing the Port Log 14-11
Configuring for syslogd 14-13
Configuring the Host 14-14
Configuring the Switch 14-14
Viewing and Saving Diagnostic Information 14-15
Setting Up Automatic Trace Dump Transfers 14-16

Chapter 15 Troubleshooting

About Troubleshooting 15-1
Most Common Problem Areas 15-2
Gathering Information for Technical Support. 15-3
Troubleshooting Questions 15-3
Analyzing Connection Problems 15-5
Restoring a Segmented Fabric 15-8

Correcting Zoning Setup Issues 15-10
Recognizing MQ-WRITE Errors 15-12
Correcting I2C Bus Errors 15-13
Correcting Device Login Issues 15-14
Identifying Media-Related Issues 15-19
Correcting Link Failures 15-21
Correcting Marginal Links 15-24
Inaccurate Information in the System Message Log 15-25
Recognizing the Port Initialization and FCP Auto Discovery Process 15-26
Port Mirroring 15-27
Supported Hardware 15-27
How Port Mirroring Works 15-28
Port Mirroring Considerations 15-30
Creating, Deleting, and Displaying Port Mirroring 15-31

Chapter 16 Administering NPIV

About NPIV 16-1
Enabling and Disabling NPIV 16-2
Configuring NPIV 16-2
Configuration Scenarios 16-3
Viewing NPIV Port Configuration Information 16-3
Displaying Login Information 16-6

Section II Licensed Features

Chapter 17 Administering Advanced Performance Monitoring

About Advanced Performance Monitoring. 17-1
Displaying and Clearing the CRC Error Count 17-3
Monitoring End-to-End Performance 17-3
Adding End-to-End Monitors 17-4
Setting a Mask for End-to-End Monitors 17-6
Deleting End-to-End Monitors 17-8
Monitoring Filter-Based Performance 17-8
Adding Standard Filter-based Monitors 17-9
Adding Custom Filter-Based Monitors. 17-10
Deleting Filter-Based Monitors 17-11
Monitoring ISL Performance 17-12
Monitoring Trunks 17-12
Displaying Monitor Counters 17-13
Clearing Monitor Counters 17-15
Saving and Restoring Monitor Configurations 17-16
Collecting Performance Data 17-17

Chapter 18 Administering Extended Fabrics

About Extended Link Buffer Allocation. 18-1
SilkWorm 3250, 3850, 3900, 24000, and 48000 (FC2-16 Port Blades) . . 18-1
SilkWorm 4100, 4900, 7500, and 48000 (FC4-16, FC4-16IP, FC4-32, FC4-48, and FR4-18i Port Blades) 18-2
Fabric Considerations 18-2
Choosing an Extended ISL Mode 18-3
Switches with Bloom ASIC 18-4
Switches with Goldeneye ASIC 18-5
Switches with Condor ASIC. 18-5
Configuring External Ports 18-7
Configuring an Extended ISL 18-10

Chapter 19 Administering ISL Trunking

About ISL Trunking 19-1
Standard Trunking Criteria 19-2
Fabric Considerations 19-3
Initializing Trunking on Ports 19-4
Monitoring Traffic 19-5
Enabling and Disabling ISL Trunking 19-6
Setting Port Speeds 19-7
Displaying Trunking Information 19-8
Trunking Over Extended Fabrics 19-9
Troubleshooting Trunking Problems 19-11
    Listing Link Characteristics 19-11
    Recognizing Buffer Underallocation 19-12

Chapter 20 Administering Advanced Zoning

About Zoning 20-1
    Zone Types 20-3
    Zone Objects 20-5
    Zone Aliases 20-5
    Zone Configurations 20-6
    Zoning Enforcement 20-6
    Rules for Configuring Zones 20-10
Creating and Managing Zone Aliases 20-11
Creating and Maintaining Zones 20-13
    Activating Default Zones 20-15
    Merging Zones 20-16
Creating and Modifying Zoning Configurations 20-19
    Maintaining Zone Objects 20-23
Managing Zoning Configurations in a Fabric 20-26
    Adding a New Switch or Fabric 20-26
    Splitting a Fabric 20-28

Using Zoning to Administer Security 20-28
Resolving Zone Conflicts 20-29

Chapter 21 Configuring and Monitoring FCIP Tunneling

Supported Platforms for FCIP 21-1
About FCIP Tunneling 21-1
FCIP Licensing 21-2
FCIP Tunneling Concepts 21-3
Port Numbering 21-5
Tunneling and IPSec 21-6
Configuring IPSec 21-8
Managing Policies 21-9
Configuring FCIP Tunnels 21-11
Enabling Persistently Disabled Ports 21-12
Defining the IP Interface of Each Virtual Port 21-13
Configuring the GbE Ports 21-13
Adding IP Routes on a GbE Port 21-14
Verifying IP Connectivity 21-15
Testing End-to-End IP Path Performance 21-15
Fastwrite and Tape Pipelining 21-16
Configuring FCIP Tunnels 21-20
Verifying the FCIP Tunnel Configuration 21-22
Checklist for Configuring FCIP Links 21-25
Troubleshooting FCIP Links 21-25
WAN Performance Analysis Tools 21-26
About the Ipperf Option 21-26
WAN Tool Performance Characteristics 21-27
Starting WAN Tool Analysis 21-28

Appendix A Configuring the PID Format

About PIDs and PID Binding A-1
Summary of PID Formats A-2
Impact of Changing the Fabric PID Format A-2
Host Reboots A-3
Static PID Mapping Errors A-3
Changes to Configuration Data A-3
Selecting a PID format A-4
Evaluating the Fabric A-6
Planning the Update Procedure A-7
Online Update A-8
Offline Update A-8
Hybrid Update A-9
Changing to Core PID Format A-9
Changing to Extended Edge PID Format A-10
Converting Port Number to Area ID A-12
Performing PID Format Changes A-14
Basic Procedure A-14
HP/UX Procedure A-15
AIX Procedure A-17
Swapping Port Area IDs A-18

Appendix B Configuring Interoperability Mode

Vendor Switch Requirements B-1
Brocade Switch Requirements B-2
Supported Brocade Features B-2
Unsupported Brocade Features B-2
Configuration Recommendations B-3
Configuration Restrictions B-3
Zoning Restrictions B-4
Zone Name Restrictions B-5
Enabling and Disabling Interoperability Mode B-5

Appendix C Understanding Legacy Password Behavior

Password Management Information C-1
Password Prompting Behaviors C-3
Password Migration During Firmware Changes C-4
Password Recovery Options C-5

Appendix D Using Remote Switch

About Remote Switch D-1
Remote Switch Capabilities D-1
Using Remote Switch with a Gateway D-2

Appendix E Zone Merging Scenarios

Index

About This Document

This document is a procedural guide to help SAN administrators configure and manage a storage area network (SAN) using the Brocade Fabric OS Command Line Interface (CLI).

How This Document Is Organized

The document is divided into two sections; the first, "Standard Features," contains the following topics:

- Chapter 1, "Introducing Fabric OS CLI Procedures," gives a brief overview of Fabric OS, discusses the differences between SilkWorm switches and directors, and explains the Fabric OS CLI Help feature.
- Chapter 2, "Performing Basic Configuration Tasks," provides typical connection and configuration procedures.
- Chapter 3, "Managing User Accounts," provides information and procedures on managing authentication and user accounts for the switch management channel.
- Chapter 4, "Configuring Standard Security Features," provides procedures for basic password and user account management.
- Chapter 5, "Configuring Advanced Security," provides information and procedures for configuring ACL policies for FC port and switch binding and managing the fabric-wide consistency policy.
- Chapter 6, "Maintaining Configurations," provides procedures for maintaining and backing up your switch configurations.
- Chapter 7, "Managing Administrative Domains," describes the concepts and provides procedures for using administrative domains.
- Chapter 8, "Installing and Maintaining Firmware," provides preparations and procedures for performing firmware downloads, as well troubleshooting information.
- Chapter 9, "Configuring SilkWorm Directors," provides information and procedures specific to SilkWorm 24000 and 48000 models. Because the SilkWorm 24000 and 48000 models have CP blades and port blades, they require procedures that are not relevant to the SilkWorm 3250, 3850, 3900, 4100, 4900 and 7500 fixed-port models.
- Chapter 10, "Routing Traffic," provides information and procedures for using switch routing features.
- Chapter 11, "Using the FC-FC Routing Service," provides information for setting up and using the FC-FC Routing Service.
- Chapter 12, "Administering FICON Fabrics," provides procedures for use of the Brocade FICON Fabrics licensed feature.
- Chapter 13, "Configuring the Distributed Management Server," describes the use of the SAN management application.

- Chapter 14, “Working With Diagnostic Features,” provides information about diagnostic and status-determining features, particularly system message logging.
- Chapter 15, “Troubleshooting,” provides problem solving information and procedures.
- Chapter 16, “Administering NPIV,” provides procedures for enabling and configuring N-Port ID Virtualization (NPIV).

The second section, “Licensed Features,” contains the following topics:

- Chapter 17, “Administering Advanced Performance Monitoring,” provides procedures for use of the Brocade Advanced Performance Monitoring licensed feature.
- Chapter 18, “Administering Extended Fabrics,” provides procedures for use of the Brocade Extended Fabrics licensed feature.
- Chapter 19, “Administering ISL Trunking,” provides procedures for use of the Brocade ISL Trunking licensed feature.
- Chapter 20, “Administering Advanced Zoning,” provides procedures for use of the Brocade Advanced Zoning licensed feature.
- Chapter 21, “Configuring and Monitoring FCIP Tunneling,” provides procedures for creating and maintaining FCIP tunnels.
- Appendix A, “Configuring the PID Format,” provides information about the various switch PID formats available and procedures for setting the PID format.
- Appendix B, “Configuring Interoperability Mode,” provides information about using SilkWorm switches with other brands of switches.
- The other appendices provide special procedures or legacy information for previous Fabric OS versions.

Supported Hardware and Software

This document is specific to Fabric OS v5.2.0 running on the following Brocade SilkWorm product models:

- Brocade SilkWorm 200E switch
- Brocade SilkWorm 3250 switch
- Brocade SilkWorm 3850 switch
- Brocade SilkWorm 3900 switch
- Brocade SilkWorm 4100 switch
- Brocade SilkWorm 4900 switch
- Brocade SilkWorm 7500 switch
- Brocade SilkWorm 24000 director
- Brocade SilkWorm 48000 director

When procedures or parts of procedures documented here apply to some switches but not to others, this guide identifies which switches are supported and which are not.

This document sometimes mentions other Fabric OS releases to highlight the changes in the latest release or to point out interoperability issues with other SilkWorm models. It also specifies when procedures or steps of procedures apply only to specific SilkWorm models.

Although many different software and hardware configurations are tested and supported by Brocade Communications Systems, Inc. for this Brocade Fabric OS release, documenting all possible configurations and scenarios is beyond the scope of this document.

What's New in This Document

The following changes have been made since this document was last released:

- Information that was added:
  - Admin Domains
  - Fastwrite and tape pipelining
  - FC-FC Routing Service features:
    - Front domain consolidation
    - FC Router port cost
    - EX_Port frame trunking
  - IPSec
  - Port mirroring
  - Security features:
    - Role-Based Access Control (RBAC)
    - Access Control List (ACL)
  - Support for the SilkWorm FC4-48 and FC4-16IP port blades
  - WAN Tools
- Information that was changed:
  - Configuring audit events
  - Configuring DHCP
  - FC-FC Routing:
    - Monitoring resources
    - McDATA SANs
  - FCIP
  - Firmware download
  - Ports on Demand
  - RADIUS server
  - Time Zone now supports Daylight Savings
  - User management
- Information that was deleted:
  - Mention of support for the SilkWorm 3014, 3016, and 4012 on Fabric OS v5.2.0
  - Information about QuickLoop

For further information, refer to the release notes.

Document Conventions

This section describes text formatting conventions and important notices formats.

Text Formatting

The narrative-text formatting conventions that are used in this document are as follows:

bold text	Identifies command names Identifies the names of user-manipulated GUI elements Identifies keywords and operands Identifies text to enter at the GUI or CLI
italic text	Provides emphasis Identifies variables Identifies paths and Internet addresses Identifies document titles
`code` text	Identifies CLI output Identifies syntax examples

For readability, command names in the narrative portions of this guide are presented in mixed lettercase: for example, **switchShow**. In actual examples, command lettercase is often all lowercase. Otherwise, this manual specifically notes those cases in which a command is case sensitive. The **ficonCupSet** and **ficonCupShow** commands are an exception to this convention.

Notes, Cautions, and Warnings

The following notices appear in this document.

Note
A note provides a tip, emphasizes important information, or provides a reference to related information.

Caution
A caution alerts you to potential damage to hardware, firmware, software, or data.

Warning
A warning alerts you to potential danger to personnel.

Key Terms

For definitions specific to Brocade and Fibre Channel, see the *Brocade Glossary.*

For definitions of SAN-specific terms, visit the Storage Networking Industry Association online dictionary at:

http://www.snia.org/education/dictionary.

Additional Information

This section lists additional Brocade and industry-specific documentation that you might find helpful.

Brocade Resources

The following related documentation is provided on the Brocade Documentation CD-ROM and on the Brocade Web site, through Brocade Connect.

Note

Go to http://www.brocade.com and click **Brocade Connect** to register at no cost for a user ID and password.

Fabric OS

- *Fabric OS Command Reference*
- *Fabric OS MIB Reference*
- *Fabric OS Message Reference*
- *Brocade Glossary*

Optional Applications

- *Fabric Watch Administrator's Guide*
- *Fabric Manager Administrator's Guide*
- *Secure Fabric OS Administrator's Guide*
- *Web Tools Administrator's Guide*
- *EZSwitchSetup Administrator's Guide*

SilkWorm 48000

- *SilkWorm 48000 Hardware Reference Manual*
- *SilkWorm 48000 QuickStart Guide*
- *FR4-18i Hardware Reference Manual*
- *FC4-16IP Hardware Reference Manual*

SilkWorm 24000

- *SilkWorm 24000 Hardware Reference Manual*
- *SilkWorm 24000 QuickStart Guide*

SilkWorm 24000/48000

- *Port Blade and Filler Panel Replacement Procedure*
- *Control Processor Blade Replacement Procedure*
- *Blower Assembly Replacement Procedure*
- *Cable Management Tray and Guide Replacement Procedure*
- *Chassis Door Replacement Procedure*

- *WWN Bezel and Card Replacement Procedure*
- *Power Supply and Filler Panel Replacement Procedure*
- *14U Rack Mount Kit Installation Procedure*
- *Mid-Mount Rack Kit Installation Procedure*

SilkWorm 7500

- *SilkWorm 7500 Hardware Reference Manual*
- *SilkWorm 7500 QuickStart Guide*
- *SilkWorm 7500 Fan Assembly Replacement Procedure*
- *SilkWorm Mid Sized Power Supply Replacement Procedure*

SilkWorm 4900

- *SilkWorm 4900 Hardware Reference Manual*
- *SilkWorm 4900 QuickStart Guide*
- *SilkWorm 4900 Fan Assembly Replacement Procedure*
- *SilkWorm Mid Sized Power Supply Replacement Procedure*

SilkWorm 4100

- *SilkWorm 4100 Hardware Reference Manual*
- *SilkWorm 4100 QuickStart Guide*

SilkWorm 3900

- *SilkWorm 3900 Hardware Reference Manual* (for v4.x software)
- *SilkWorm 3900 QuickStart Guide* (for v4.x software)
- *SilkWorm 3900 Fan Assembly Replacement Procedure*
- *SilkWorm 3900 Motherboard Assembly Replacement Procedure*
- *SilkWorm 3900 Power Supply Replacement Procedure*

SilkWorm 3250/3850

- *SilkWorm 3250/3850 Hardware Reference Manual* (for v4.x software)
- *SilkWorm 3250/3850 QuickStart Guide* (for v4.x software)

SilkWorm 200E

- *SilkWorm 200E Hardware Reference Manual* (for v5.x software)

For practical discussions about SAN design, implementation, and maintenance, you can obtain *Building SANs with Brocade Fabric Switches* through:

http://www.amazon.com

For information about how to use Fabric OS features in a SAN solution, refer to the *Principles of SAN Design* (ISBN 0-7414-2824-5), available from the SAN Administrator's Bookshelf. You can order the book and view a sneak preview at:

http://www.brocade.com/products/sanadmin_bookshelf/index.jsp

For additional Brocade documentation, visit the Brocade SAN Info Center and click the Resource Library location:

http://www.brocade.com

Release Notes are available on the Brocade Connect Web site and are also bundled with the Fabric OS firmware.

Optional Brocade Features

Optional Brocade features include:

Advanced Performance Monitoring
: Enables more effective end-to-end SAN performance analysis to enhance performance tuning, increase productivity, optimize resource utilization, and reduce costs.

Extended Fabrics
: Supports the reliable, high-speed connectivity of SilkWorm switches over dark fiber or Dense Wave Division Multiplexing (DWDM) equipment at distances up to 500 kilometers to enhance business continuance operations.

Fabric Watch
: Continuously monitors SAN fabrics for potential faults based on thresholds set for a variety of SAN fabric elements and events—automatically alerting administrators to potential problems before they become costly failures.

ISL Trunking
: Optimizes the performance and availability of SAN fabrics while simplifying ISL management. Two 4 Gbit/sec SilkWorm switches can automatically group up to eight ISLs into a single logical “trunk” with a total throughput of up to 32 Gbit/sec.

Advanced Zoning
: Automatically groups SAN fabric-connected devices into logical zones that restrict access to “member” devices in the zone. Advanced Zoning uses hardware enforcement at both the port and WWN level to provide more robust data protection.

Secure Fabric OS
: Provides a comprehensive security solutions to help protect mission-critical data. Key features include centralized policy-based security management, encryption of management data, and authentication to create a fabric-wide trusted environment with control over all levels of fabric access and communication.

FICON CUP
: Enables IBM host-based management programs to manage FICON fabric switches in-band by sending commands to the Fabric OS emulated control device.

Other Industry Resources

For additional resource information, visit the Technical Committee T11 Web site. This Web site provides interface standards for high-performance and mass storage applications for Fibre Channel, storage management, as well as other applications:

http://www.t11.org

For information about the Fibre Channel industry, visit the Fibre Channel Industry Association Web site:

http://www.fibrechannel.org

Getting Technical Help

Contact your switch support supplier for hardware, firmware, and software support, including product repairs and part ordering. To expedite your call, have the following information available:

1. **General Information**
   - Technical Support contract number, if applicable
   - Switch model
   - Switch operating system version
   - Error numbers and messages received
   - **supportSave** command output
   - Detailed description of the problem and specific questions
   - Description of any troubleshooting steps already performed and results
   - Serial console and telnet session logs
   - syslog message logs
2. **Switch Serial Number**

   The switch serial number and corresponding bar code are provided on the serial number label, as illustrated below.

   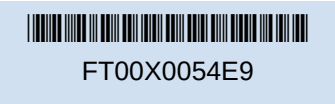

   The serial number label is located as follows:

   - *SilkWorm 3014*—On the top of the chassis, under the insertion arm
   - *SilkWorm 3016 and 4012*—On the bottom of the switch module
   - *SilkWorm 200E, 3200, 3250 and 3850*—On the bottom of the chassis
   - *SilkWorm 3800 and 3900*— Nonport side of the chassis
   - *SilkWorm 4100, 4900, and 7500*—On the switch ID pull-out tab located inside the chassis on the port side on the left
   - *SilkWorm 12000, 24000, and 48000*—Inside the chassis next to the power supply bays
   - *SilkWorm Multiprotocol Router Model AP7420*—On the bottom of the chassis and on the back of the chassis.
3. **World Wide Name (WWN)**
   - *SilkWorm 200E, 3014, 3016, 3250, 3800, 3850, 3900, 4012, 4100, 4900, 7500 switches and SilkWorm 24000, and 48000 directors:* Provide the license ID. Use the **licenseIdShow** command to display the license ID.
   - *SilkWorm Multiprotocol Router Model AP7420:* Provide the switch WWN. Use the **switchShow** command to display the switch WWN.
   - *All other SilkWorm switches:* Provide the switch WWN. Use the **wwn** command to display the switch WWN.

Document Feedback

Because quality is our first concern at Brocade, we have made every effort to ensure the accuracy and completeness of this document. However, if you find an error or an omission, or you think that a topic needs further development, we want to hear from you. Forward your feedback to:

documentation@brocade.com

Provide the title and version number and as much detail as possible about your issue, including the topic heading and page number and your suggestions for improvement.

Section

I

Standard Features

This section describes standard Brocade Fabric OS features, and includes the following chapters:

- Chapter 1, “Introducing Fabric OS CLI Procedures”
- Chapter 2, “Performing Basic Configuration Tasks”
- Chapter 3, “Managing User Accounts”
- Chapter 4, “Configuring Standard Security Features”
- Chapter 5, “Configuring Advanced Security”
- Chapter 6, “Maintaining Configurations”
- Chapter 7, “Managing Administrative Domains”
- Chapter 8, “Installing and Maintaining Firmware”
- Chapter 9, “Configuring SilkWorm Directors”
- Chapter 10, “Routing Traffic”
- Chapter 11, “Using the FC-FC Routing Service”
- Chapter 12, “Administering FICON Fabrics”
- Chapter 13, “Configuring the Distributed Management Server”
- Chapter 14, “Working With Diagnostic Features”
- Chapter 15, “Troubleshooting”
- Chapter 16, “Administering NPIV”

Chapter

1

Introducing Fabric OS CLI Procedures

This guide contains procedures for configuring and managing a Brocade® SilkWorm® Storage Area Network (SAN) using the Fabric OS Command Line Interface (CLI).

- "About Procedural Differences" on page 1-1
- "Scope and References" on page 1-2
- "About the CLI" on page 1-2
- "Help Information" on page 1-3

The guide applies to the following Brocade product models:

- SilkWorm switches: 200E, 3250, 3850, 3900, 4100, 4900, and 7500

  These SilkWorm platforms contain a fixed number of ports (they are "fixed-port switches"). The SilkWorm 200E, 4100, and 4900 and models allow you to license and activate extra fixed ports with the Ports on Demand (POD) feature.

- SilkWorm directors: 24000 and 48000

  These SilkWorm platforms can contain a variable number of ports, which you install by plugging port blades into the director chassis. The SilkWorm 48000 can have up to 256 ports, and the SilkWorm 24000 can have up to 128 ports.

About Procedural Differences

As a result of the differences between fixed-port and variable-port devices, procedures sometimes differ among SilkWorm models. As new SilkWorm models are introduced, new features sometimes apply only to those models.

When procedures or parts of procedures apply to some models but not others, this guide identifies the specifics for each model. For example, a number of procedures that apply only to variable-port devices are found in "Configuring SilkWorm Directors" on page 9-1. Procedures that apply only to the SilkWorm 4100 model, for example, are labeled as such.

Note

When command examples in this guide show user input enclosed in quotation marks, the quotation marks are required.

Scope and References

Although many different software and hardware configurations are tested and supported by Brocade Communications Systems, Inc., documenting all possible configurations and scenarios is beyond the scope of this document. In some cases, earlier releases are highlighted to present considerations for interoperating with them.

The hardware reference manuals for SilkWorm products describe how to power up devices and set their IP addresses. After the IP address is set, you can use the CLI procedures contained in this guide.

This guide provides only the level of detail required to perform the procedures. If you need more information about the commands used in the procedures, see online help or to the *Fabric OS Command Reference*.

There are several access methods that you can use to configure a switch. These are listed with their respective documents:

- Command line interface (CLI)
  - A telnet session into logical switches
  - A telnet session into active and standby CPs for director class switches
  - A serial console, including active and standby CPs for director class switches
  - An optional modem, which behaves like a serial console port

  For CLI details, see the *Fabric OS Command Reference*.
- Web Tools

  For Web Tools procedures, see the *Web Tools Administrator's Guide*.
- Fabric Manager

  For Fabric Manager procedures, see the *Fabric Manager Administrator's Guide*.
- A third-party application using the API

  For third-party application procedures, refer to the third-party API documentation.

About the CLI

Fabric OS CLI is the complete fabric management tool for Brocade SANs that enables you to:

- Access the full range of Fabric OS features, based on license keys.
- Configure, monitor, dynamically provision, and manage every aspect of the SAN.
- Configure and manage the Brocade fabric on multiple efficient levels.
- Identify, isolate, and manage SAN events across every switch in the fabric.
- Manage switch licenses.
- Perform fabric stamping.

To manage a switch using telnet, SSH2 session into SNMP, or Web Tools, the switch must be connected to a network through the switch Ethernet port (out of band) or from the Fibre Channel (in band). The switch must be configured with an IP address to allow for the network connection. Refer to the hardware manual for your specific switch for information on physically connecting to the switch.

You can access switches from different connections, such as Web Tools, CLI, and API. When these connections are simultaneous, changes from one connection might not be updated to the other, and some modifications might be lost. When simultaneous connections are used, make sure that you do not overwrite the work of another connection.

In a mixed fabric containing switches running various Fabric OS versions, you should use the latest-model switches running the most recent release for the primary management tasks. The principal management access should be set to the core switches in the fabric. For example, to run Secure Fabric OS®, use the latest-model switch as the primary FCS, the location to perform zoning tasks, and the time server.

A number of management tasks are designed to make fabric-level changes; for example, zoning commands make changes that affect the entire fabric. When executing fabric-wide configuration tasks, allow time for the changes to propagate across the fabric before executing any subsequent tasks. For a large fabric, it might take a few minutes.

Help Information

Each Fabric OS command provides Help information that explains the command function, its possible operands, its level in the command hierarchy, and additional pertinent information.

Displaying Command Help

1. Connect to the switch and log in as admin.
2. To display a list of all command help topics for a given login level, enter the **help** command with no arguments.

   For example, if you are logged in as *user* and enter the **help** command, a list of all *user-level commands* that can be executed will be listed. The same rule applies to the *admin* role. In addition, any user-configured command that uses a *switchAdmin* role will also see commands available to users with the *switchAdmin* role.

3. To display help for a specific command for a given login level, enter **help** <***command***>, where *command* is the name of the command for which you need information.

```
switch:admin> help timeout
Administrative Commands                                  timeout(1m)

NAME
     timeout  -  Sets  or  displays the timeout value for a login session.

SYNOPSIS
     timeout [timeval]

AVAILABILITY
     admin       (set)

     all users   (display)

DESCRIPTION
     Use this command without any operands to  display  (in  min-utes)  the
current      timeout value after which idle logins are automatically terminated.

     Use this command with the timeval operand to set  the  login timeout  value
to      the specified interval.  A value of 0 will disable the timeout
functionality of      login sessions.
(output truncated)
```

Displaying Additional Help Topics

The following commands provide help files for specific topics:

diagHelp	Diagnostic help information
ficonHelp	FICON help information
fwHelp	Fabric Watch help information
iscsiHelp	iSCSI help informations
licenseHelp	License help information
perfHelp	Performance Monitoring help information
routeHelp	Routing help information
trackChangesHelp	Track Changes help information
zoneHelp	Zoning help information

Chapter

2

Performing Basic Configuration Tasks

This chapter contains the following procedures for performing basic switch configuration tasks using the Fabric OS command line interface (CLI).

- "Connecting to the Command Line Interface" on page 2-2
- "Setting the Default Account Passwords" on page 2-4
- "Setting Static Ethernet Addresses" on page 2-7
- "Setting the Date and Time" on page 2-9
- "Maintaining Licensed Features" on page 2-14
- "Customizing Switch Names" on page 2-16
- "Customizing the Chassis Name" on page 2-17
- "Working With Domain IDs" on page 2-17
- "Activating Ports on Demand" on page 2-19
- "Disabling and Enabling Switches" on page 2-20
- "Disabling and Enabling Ports" on page 2-21
- "Making Basic Connections" on page 2-22
- "Linking Through a Gateway" on page 2-23
- "Checking Status" on page 2-24
- "Tracking and Controlling Switch Changes" on page 2-25
- "Configuring the Audit Log" on page 2-29
- "Shutting Down Switches and Directors" on page 2-32

And this topic:

- "High Availability of Daemon Processes" on page 2-33

Ideally, perform the initial configuration of a switch prior to introducing the switch into the fabric or during a scheduled maintenance window to minimize fabric disruption.

Connecting to the Command Line Interface

Connect to the Fabric OS command line interface (CLI) either through a telnet or SSH connection or through a console session on the serial port.

Using Telnet or SSH session

Connect to the Fabric OS CLI of a switch that has a configured network interface using a telnet or SSH session. The switch must also be physically connected to the network. If the switch network interface is not configured or the switch has been disconnected from the network, use a console session on the serial port as described in the next section.

Note

To automatically configure the network interface on a DHCP-enabled switch, plug the switch into the network and power it on. The DHCP client automatically gets the IP and gateway addresses from the DHCP server. The DHCP must be on the same subnet as the switch. See "Configuring DHCP" on page 2-7 for more details.

Observe these rules for telnet connections:

- Never change the IP address of the switch while two telnet sessions are active; if you do, your next attempt to log in fails. To recover, gain access to the switch by one of these methods:
  - You can use Web Tools and perform a fast boot. When the switch comes up, the telnet quota is cleared. (For instructions on performing a fast boot with Web Tools, see the *Web Tools Administrator's Guide*.)
  - If you have the required privileges, you can connect through the serial port, log in as root, and use operating system commands to identify and kill the telnet processes without disrupting the fabric.
- For admin level accounts, Fabric OS limits the number of simultaneous telnet sessions per switch to two. For more details on session limits, see "Configuring the Telnet Interface" on page 4-4 and "Managing User Accounts" on page 3-1.

To connect using telnet

1. Verify that the switch's network interface is configured and that it is connected to the (IP) network through the RJ-45 Ethernet port.

   Switches in the fabric that are not connected via Ethernet can be managed through switches that are using IP over Fibre Channel. The embedded port must have an assigned IP address.

2. Open a telnet connection using the IP address of the logical switch to which you want to connect.

   If you telnet to the active CP or log in to the active CP console, you are prompted for the switch number when the platform is set up in dual (or multiple) switch mode. For example, a SilkWorm 24000 will not prompt you if you are using configuration option 1, but will prompt you if you have used configuration options 2-4. See Chapter 9, "Configuring SilkWorm Directors" for details about director configuration options.

   The SilkWorm 24000 (configured with two domains) has two logical switches (**sw0** and **sw1**).

   The login prompt is displayed when the telnet connection finds the switch in the network.

3. Enter the account ID at the login prompt.

   See "Setting the Default Account Passwords" on page 2-4 for instructions on how to log in for the first time.

4. Enter the password.

   If you have not changed the system passwords from the default, you are prompted to change them. Enter the new system passwords, or press **Ctrl-c** to skip the password prompts. See "Changing Default Account Passwords at Login" on page 2-4.

5. Verify that the login was successful.

   The prompt displays the switch name and user ID to which you are connected.

```
login: admin
password: xxxxxxx
switch:admin>
```

Using a console session on the serial port

Note the following behaviors for serial connections:

- Some procedures require that you connect through the serial port; for example, setting the IP address or setting the boot PROM password.
- If secure mode is enabled, connect through the serial port of the primary FCS switch.
- **SilkWorm 24000 and 48000**: You can connect to CP0 or CP1 using either of the two serial ports.

To connect through the serial port

1. Connect the serial cable to the serial port on the switch and to an RS-232 serial port on the workstation.

   If the serial port on the workstation is RJ-45 instead of RS-232, remove the adapter on the end of the serial cable and insert the exposed RJ-45 connector into the RJ-45 serial port on the workstation.

2. Open a terminal emulator application (such as HyperTerminal on a PC, or TERM, TIP, or Kermit in a UNIX environment), and configure the application as follows:

   - In a Windows environment:

Parameter	Value
Bits per second	9600
Databits	8
Parity	None
Stop bits	1
Flow control	None

   - In a UNIX environment, enter the following string at the prompt: **tip /dev/ttyb -9600**.

     If ttyb is already in use, you can use ttya (enter **tip /dev/ttya -9600**).

Setting the Default Account Passwords

The switch automatically prompts you to change the default account passwords after logging in for the first time. If you do not change the passwords, the switch prompts you after each subsequent login until all the default passwords have been changed.

Note

The default account passwords can be changed from their original value only when prompted immediately following the login only; the passwords cannot be changed using the `passwd` command later in the session. If you skip the prompt, but later decide to change the passwords, log out and back in.

The default accounts on the switch are admin, user, root, and factory. Use the default administrative account as shown in Table 2-1, to log in to the switch for the first time and to perform the basic configuration tasks described in this chapter.

Every logical switch (domain) has a set of default accounts. The root and factory default accounts are reserved for development and manufacturing. The user account is primarily used for system monitoring. For more information on default accounts, see "About the Default Accounts" on page 3-7.

The following table describes the default administrative accounts for switches by model number:

Table 2-1 Default Administrative Account Names and Passwords

Model	Administrative account	Password
SilkWorm 200E, 3250, 3850, 3900, 4100, 4900, and **7500 switches**	admin	password
SilkWorm 24000[a] and **48000 directors**	admin	password

a. **SilkWorm 24000 configured with two domains**: Each of the two logical switches has its own set of default accounts. The default account names and passwords are the same for both of the logical switches.

Changing Default Account Passwords at Login

The "change default account passwords" prompt accepts a maximum of eight characters. Any characters beyond the eighth character are ignored. Only the default password is subject to the eight character limit. User-defined passwords can have 8 to 40 characters. They must begin with an alphabetic character and can include numeric characters, the dot (.), and the underscore (_). They are case-sensitive, and they are not displayed when you enter them on the command line.

Record the passwords exactly as entered and store them in a secure place, because recovering passwords requires significant effort and fabric downtime. Although the root and factory accounts are not meant for general use, change their passwords if prompted to do so and save the passwords in case they are needed for recovery purposes.

Note

See "Connecting to the Command Line Interface" on page 2-2 for details on how to establish a connection to the switch.

To change the default account passwords at login

1. Connect to the switch and log in using the default administrative account.
2. At each of the “Enter new password” prompts, either enter a new password or skip the prompt.

   To skip a prompt press **Enter**. Bypass the remaining prompts by pressing **Ctrl-c**.

```
login: admin
Password:
Please change your passwords now.
Use Control-C to exit or press 'Enter' key to proceed.
for user - root
Changing password for root
Enter new password: ********
Password changed.
Saving password to stable storage.
Password saved to stable storage successfully.
Please change your passwords now.
for user - factory
Changing password for factory
Enter new password: ********
Password changed.
Saving password to stable storage.
Password saved to stable storage successfully.
Please change your passwords now.
for user - admin
Changing password for admin
Enter new password: ********
Password changed.
Saving password to stable storage.
Password saved to stable storage successfully.
Please change your passwords now.
for user - user
Changing password for user
Enter new password: ********
Password changed.
Saving password to stable storage.
Password saved to stable storage successfully.
switch:admin>
```

Configuring the Ethernet Interface

Fabric OS v5.2.0 introduces the use of DHCP (Dynamic Host Configuration Protocol) for the Ethernet network interface configuration. The Ethernet (network) interface provides management access, including direct access to the Fabric OS CLI and allows other tools such as Web Tools to interact with the switch.

You can continue to use static Ethernet addressing system or allow the DHCP client to automatically acquire Ethernet addresses. Configure the Ethernet interface IP, subnet mask, and gateway addresses in one of the following manners:

- "Setting Static Ethernet Addresses" on page 2-7
- "Configuring DHCP" on page 2-7

Note

When you change the Ethernet interface settings, open connections such as SSH or telnet may be dropped. Reconnect using the new Ethernet IP information or change the Ethernet settings using a console session through the serial port to maintain your session through the change. You must connect through the serial port to set the Ethernet IP address if an the Ethernet network interface is not configured already. See "To connect through the serial port" on page 2-3 for details.

Displaying the Network Interface Settings

If an IP address has not been assigned to the network interface (Ethernet), you must connect to the Fabric OS CLI using a console session on the serial port, see "Using a console session on the serial port" on page 2-3. Otherwise, connect using SSH.

To display the network interface settings

1. Connect to the switch and log in as admin.
2. Enter the **ipAddrShow** command.

```
FD21:admin> ipaddrshow

SWITCH
Ethernet IP Address: 192.168.78.158
Ethernet Subnetmask: 255.255.255.0
Fibre Channel IP Address: 220.220.220.23
Fibre Channel Subnetmask: 255.255.255.0
Gateway Address: 192.168.78.1
DHCP: On
```

If the Ethernet IP address, subnet mask, and gateway address are displayed, then the network interface is configured. Verify that the information is correct. If DHCP is enabled, the network interface information was acquired from the DHCP server.

Setting Static Ethernet Addresses

Use static Ethernet network interface addresses on SilkWorm 24000 and 48000 Directors, and in environments where DHCP service is not available. To use static addresses for the Ethernet interface, you must first disable DHCP. You may enter static Ethernet information and disable DHCP at the same time.

If you choose not use DHCP or to specify an IP address for your switch Ethernet interface, you can do so by entering none or 0.0.0.0 in the Ethernet IP address field.

Caution
The use of IP address 0.0.0.0 is not supported in pre-Fabric OS v5.2.0 fabrics.

Fabric OS beginning with v2.6.0, v3.1.0, and v4.0.0 supports Classless Inter-Domain Routing (CIDR).

To set static addresses for the Ethernet network interface

1. Connect to the switch and log in as admin.
2. Enter the following command:

```
ipaddrset [-cp number] [-sw number]
```

```
switch:admin> ipaddrset -sw 1
Ethernet IP Address [192.168.166.148]:
Ethernet Subnetmask [255.255.255.0]:
Fibre Channel IP Address [none]:
Fibre Channel Subnetmask [none]:
Committing configuration...Done.
OK.
switch:root>
```

Enter the network information in dotted quad format for Ethernet IP address, Ethernet Subnetmask, and Gateway Address at the prompts. Skip Fibre Channel prompts by pressing enter. Disable DHCP by entering Off.

Configuring DHCP

By default, some SilkWorm switches have DHCP enabled, see the Fabric OS Release Notes for a complete list of switches. SilkWorm 24000 and 48000 Directors do not support DHCP.

The Fabric OS DHCP client supports the following parameters:

- External Ethernet port IP addresses and subnet masks
- Default gateway IP address

The DHCP client uses a DHCP vendor class identifier that allows DHCP servers to determine that the Discovers and Requests are coming from a Brocade switch. The vendor class identifier is the string “BROCADE” followed by the SWBD model number of the platform. For example, the vendor class identifier for a request from a SilkWorm 3900 is “BROCADESWBD12.”

Note

The client conforms to the latest IETF Draft Standard RFCs for IPv4 DHCP.

Enabling DHCP

Plug DHCP enabled switches in to the network, power on the switch, and the switch automatically obtains the Ethernet IP address, Ethernet subnet mask, and default gateway address from the DHCP server. The DHCP client can only connect to a DHCP server on the same subnet as the switch. Do not enable DHCP, if the DHCP server is not on the same subnet as the switch.

Enabling DHCP after the Ethernet information has been configured releases the current Ethernet network interface settings, including Ethernet IP, Ethernet Subnetmask, and Gateway. The Fibre Channel (FC) IP address and subnet mask is static and is not affected by DHCP, see “Setting Static Ethernet Addresses” on page 2-7 for instructions on setting the FC IP address.

To enable DHCP

1. Connect to the switch and log in as admin.
2. Enter the following command:

```
switch:admin> ipaddrset
Ethernet IP Address [192.168.74.102]:
Ethernet Subnetmask [255.255.255.0]:
Fibre Channel IP Address [220.220.220.2]:
Fibre Channel Subnetmask [255.255.0.0]:
Gateway IP Address [192.168.74.1]:
DHCP [Off]:on
```

Skip the Ethernet IP address, Ethernet Subnetmask, Fibre Channel IP address and subnet mask prompts by pressing enter. For DHCP[Off], enable it by entering On.

Disabling DHCP

When you disable DHCP, enter the static Ethernet IP address and subnet mask of the switch and default gateway address. Otherwise, the Ethernet settings may conflict with other addresses assigned by the DHCP server on the network.

To disable DHCP

1. Connect to the switch and log in as admin.
2. Enter the following command:

```
switch:admin> ipaddrset
Ethernet IP Address [192.168.74.102]:
Ethernet Subnetmask [255.255.255.0]:
Fibre Channel IP Address [220.220.220.2]:
Fibre Channel Subnetmask [255.255.0.0]:
Gateway IP Address [192.168.74.1]:
DHCP [Off]:off
```

Enter the network information in dotted quad format for Ethernet IP address, Ethernet Subnetmask, and Gateway Address at the prompts. If a static Ethernet address is not available when you disable DHCP, enter 0.0.0.0 at the Ethernet IP address prompt. Skip Fibre Channel prompts by pressing enter. Disable DHCP by entering Off.

Setting the Date and Time

Switches maintain the current date and time inside a battery-backed real-time clock (RTC) circuit. Date and time are used for logging events. Switch operation does not depend on the date and time; a switch with an incorrect date and time value still functions properly. However, because the date and time are used for logging, error detection, and troubleshooting, you should set them correctly.

Authorization access to set or change date and time for a switch is role-based.

Note

If secure mode is not enabled, a change in date or time to one switch is forwarded to the principal switch and distributed to the fabric. If secure mode is enabled, date or time changes can be made only on the primary FCS switch and distributed to the fabric.

To set the date and time

1. Connect to the switch and log in as admin.
2. Enter the **date** command at the command line, using the following syntax:

```
date "mmddHHMMyy"
```

The values represent the following:

- mm is the month; valid values are 01 through 12.
- dd is the date; valid values are 01 through 31.
- HH is the hour; valid values are 00 through 23.
- MM is minutes; valid values are 00 through 59.
- yy is the year; valid values are 00 through 99 (values greater than 69 are interpreted as 1970 through 1999, and values less than 70 are interpreted as 2000-2069).

```
switch:admin> date
Fri Jan 29 17:01:48 UTC 2004
switch:admin> date "0227123004"
Thu Feb 27 12:30:00 UTC 2004
switch:admin>
```

For details about how to change time zones, see **tsTimeZone** command in the *Fabric OS Command Reference*.

Setting Time Zones

Fabric OS v5.2.0 provides the capability to set the time zone for a switch by name. You can specify the setting using Country and City or PST parameters. Switch operation does not depend upon a date and time setting. However, having accurate time setting is needed for accurate logging and audit tracking.

If the time zone is not set with the new options, the switch retains the offset time zone settings. The **tsTime Zone** command includes an option to revert to the prior time zone format. For more information about this **--old** option, see the *Fabric OS Command Reference*.

Note

If you are downgrading from Fabric OS v5.2.0 or retaining the offset format, see prior versions of the *Fabric OS Administrator's Guide* for the detailed information about setting time zones using the offset format. See "About the Firmware Download Process" on page 8-2 for time zone downgrading considerations.

You can set the time zone for a switch using the **tsTimeZone** command. The **tsTimeZone** command allows you to:

- Display all of the time zones supported in the firmware
- Set the time zone based on a Country and City combination or based on a time zone ID such as PST

See the **tsTimeZone** command in the *Fabric OS Command Reference* for more detailed information about the command parameters.

The time zone setting has the following characteristics:

- Users can view the time zone settings. However, only those with administrative permissions can set the time zones.
- The **tsTimeZone** setting supports daylight savings time. It automatically adjusts for Daylight Savings Time.
- Changing the time zone on a switch updates the local time zone setup and is reflected in local time calculations.

- All switches are by default in the GMT time zone (0,0). If all switches in a fabric are in one time zone, it is possible for you to keep the time zone setup at the default setting.
- System services started at the switch boot up reflect the time zone changes only after the next reboot.
- Time zone settings are persisted across failover for high availability.

  Setting the time zone on a dual domain chassis has the following characteristics:

  - Updating the time zone on any switch, updates the entire chassis.
  - The time zone of the entire chassis is the time zone of the switch 0.
  - For dual domain Directors (SilkWorm 24000), both switches in the same chassis will be in the same time zone. Dual Domain chassis do not support different time zones on each domain.

The following procedure describes how to set the time zone for a switch. You must perform the procedure on *all* switches for which the time zone needs to be set. However, you only need to set the time zone once on each switch, because the value is written to nonvolatile memory.

To set the time zone

1. Connect to the switch and log in.
2. Enter the **tsTimeZone** command as follows:

```
switch:admin> tstimezone [--interactive]/ [, timezone_fmt]
```

   - Use **tsTimeZone** with no parameters to display the current time zone setting
   - Use **--interactive** to list all of the timezones supported by the firmware.
   - Use **timezone_fmt** to set the time zone by Country/City or by time zone ID, such as PST.

The following example shows how to display the current time zone setup and how to change the time zone to US/Central.

```
switch:admin> tstimezone
Time Zone : US/Pacific
switch:admin> tstimezone US/Central
switch:admin> tstimezone
Time Zone : US/Central
```

The following procedure illustrates how to set the current time zone using interactive mode to Pacific Standard Time.

To set the time zone interactively

1. Type the **tsTimeZone** command as follows:

```
switch:admin> tstimezone --interactive
```

2. You are first prompted to select a general location.

```
Please identify a location so that time zone rules can be set correctly.
Please select a continent or ocean.
Africa
Americas
Antarctica
Arctic Ocean
Asia
Atlantic Ocean
Australia
Europe
Indian Ocean
Pacific Ocean
none - I want to specify the time zone using the Posix TZ format.
Enter number or control-D to quit ?10
Local time is now:      Thu May 11 07:39:37 PDT 2006.
Universal Time is now:  Thu May 11 14:39:37 UTC 2006.
Is the above information OK?
Yes
No
Enter number or control-D to quit ?1
```

3. You are then prompted to select a country location.

```
Please select a country.
1) Chile                      15) Northern Mariana Islands
2) Cook Islands               16) Palau
3) Ecuador                    17) Papua New Guinea
4) Fiji                       18) Pitcairn
5) French Polynesia           19) Samoa (American)
6) Guam                       20) Samoa (western)
7) Kiribati                   21) Solomon Islands
8) Marshall Islands           22) Tokelau
9) Micronesia                 23) Tonga
10) Nauru                      24) Tuvalu
11) New Caledonia              25) US minor outlying islands
12) New Zealand                26) United States
13) Niue                       27) Vanuatu
14) Norfolk Island             28) Wallis & Futuna
Enter number or control-D to quit ?26
```

4. You are finally prompted to specify the time zone region.

```
Please select one of the following time zone regions.
1) Eastern Time
2) Eastern Time - Michigan - most locations
3) Eastern Time - Kentucky - Louisville area
4) Eastern Time - Kentucky - Wayne County
5) Eastern Time - Indiana - most locations
6) Eastern Time - Indiana - Crawford County
7) Eastern Time - Indiana - Starke County
8) Eastern Time - Indiana - Switzerland County
9) Central Time
10) Central Time - Indiana - Daviess, Dubois, Knox, Martin, Perry & Pulaski
11) Central Time - Indiana - Pike County
12) Central Time - Michigan - Dickinson, Gogebic, Iron & Menominee Counties
13) Central Time - North Dakota - Oliver County
14) Mountain Time
15) Mountain Time - south Idaho & east Oregon
16) Mountain Time - Navajo
17) Mountain Standard Time - Arizona
18) Pacific Time
19) Alaska Time
20) Alaska Time - Alaska panhandle
21) Alaska Time - Alaska panhandle neck
22) Alaska Time - west Alaska
23) Aleutian Islands
24) Hawaii
Enter number or control-D to quit ?18

The following information has been given:

United States
Pacific Time

Therefore TZ='America/Los_Angeles' will be used.
Local time is now:      Thu May 11 07:39:37 PDT 2006.
Universal Time is now:  Thu May 11 14:39:37 UTC 2006.
Is the above information OK?
Yes
No
Enter number or control-D to quit ?1
```

Synchronizing Local Time

You can synchronize the local time of the Principal or Primary Fabric Configuration Server (FCS) switch to an external NTP server. It is suggested that you synchronize time with an external NTP server, as described on page 2-13.

To synchronize local time with an external source

1. Connect to the switch and log in with an account assigned to the admin role.
2. Enter the **tsClockServer** command:

```
switch:admin> tsclockserver ipaddr
```

where *ipaddr* is the IP address of the NTP server, which the switch must be able to access. This operand is optional; by default this value is LOCL, which uses the local clock of the principal or primary switch as the clock server.

```
switch:admin> tsclockserver
LOCL
switch:admin> tsclockserver "132.163.135.131"
switch:admin> tsclockserver
132.163.135.131
switch:admin>
```

Maintaining Licensed Features

Feature licenses might be part of the licensed paperpack supplied with your switch software; if not, you can purchase licenses separately from your switch vendor, who will provide you with keys to unlock the features. License keys are provided on a per-chassis basis, so for products that support multiple logical switches (domains), a license key applies to all domains within the chassis.

To unlock a licensed feature, you can either use the license key provided in the paperpack supplied with switch software, or use the following procedure to generate a license key at the Brocade Web site (*http://www.brocade.com*).

Note

You need the following items for each chassis to be licensed:

- Transaction key in the paperpack document supplied with the switch software. Or when you purchased a license, your switch vendor might have given you a transaction key to be used for generating a software license key.
- License ID. To see a switch license ID, use the **licenseIdShow** command.

To unlock a licensed feature

1. If you already have a license key, go to step 6.

   If you do not have a license key and are using a transaction key, launch an Internet browser and go to the Brocade Web site at *http://www.brocade.com*.

2. Select **Software License Keys** from the **Products** pull-down menu.

   The Software License Keys instruction page appears.

3. Enter the requested information in the required fields and click **Next**.

   A verification screen appears.

4. Verify that the information appears correctly.

   Click **Submit** if the information displayed is correct. If the information is incorrect, click **Previous** and correct the information.

5. After the information is corrected, click **Submit**.

   An information screen displays the license keys and you will receive an e-mail with the software license key(s) and installation instructions.

6. Activate and verify the license as follows:
   a. Connect to the switch and log in as admin.
   b. Activate the license using the **licenseAdd** command.

```
switch:admin> licenseadd "key"
```

   The transaction key is case sensitive, so it must be entered exactly as it appears in the paperpack. To lessen the chance of error, copy and paste the transaction key. The quotation marks are optional.

   For SilkWorm 24000 and 48000 models, licenses are effective on both CP blades and on all logical switches, but are valid only when the CP blade is inserted into a chassis that has an appropriate license ID stored in the WWN card. If a CP is moved from one chassis to another, the license works in the new chassis only if the WWN card is the same in the new chassis. Otherwise, you will need to transfer licenses from the old WWN to the new WWN.

   For example, if you swap one CP blade at a time, or replace a single CP blade, then the existing CP blade (the active CP blade) propagates the licenses to the new CP blade.

   If you move a standby CP from one chassis to another, then the active CP will propagate its configuration (including license keys).

   c. Verify that the license was added by entering the **licenseShow** command. The licensed features currently installed on the switch are listed. If the feature is not listed, enter the **licenseAdd** command again.
   d. Some features may require additional configuration, or you might need to disable and reenable the switch to make them operational; see the feature documentation for details.

```
switch:admin> licenseshow
RzdeSee9wVlfTu:
    Web license
    Zoning license
    SES license
    Fabric license
    Remote Switch license
    Extended Fabric license
    Fabric Watch license
    Performance Monitor license
    Trunking license
    Security license
    4 Domain Fabric license
    FICON_CUP license
    N_Port ID Virtualization license
    FCIP license
    Ports on Demand license - additional 16 port upgrade
    2 Domain Fabric license
    Ports on Demand license - additional 16 port upgrade
```

To remove a licensed feature

1. Connect to the switch and log in as admin.
2. Enter the **licenseShow** command to display the active licenses.

3. Remove the license key using the **licenseRemove** command.

```
switch:admin> licenseremove "key"
```

The license key is case sensitive and must be entered exactly as given. The quotation marks are optional. After removing a license key, the optionally licensed feature is disabled when the switch is rebooted or when a switch disable or enable is performed.

4. Enter the **licenseShow** command to verify that the license is disabled.

```
switch:admin> licenseshow
bQebzbRdScRfc0iK:
    Web license
    Zoning license
switch:admin> licenseremove "bQebzbRdScRfc0iK"
removing license key "bQebzbRdScRfc0iK"
switch:admin>
```

After a reboot (or **switchDisable** and **switchEnable**), only the remaining licenses appear:

```
switch:admin> licenseshow
SybbzQQ9edTzcc0X:
    Fabric license
switch:admin>
```

If there are no license keys, **licenseShow** displays "No licenses."

Customizing Switch Names

Switches can be identified by IP address, Domain ID, World Wide Name (WWN), or by customized switch names that are unique and meaningful.

Version 4.0.0 (and later) switch names can be from 1 to 15 characters long, must begin with a letter, and can contain letters, numbers, or the underscore character. It is not necessary to use quotation marks.

The default names are:

- **SilkWorm 200E, 3250, 3850, 3900, 4100, 4900, and 7500 switches**: swd77
- **SilkWorm 24000 and 48000**: swd77

Note

Changing the switch name causes a domain address format RSCN (registered state change notification) to be issued, and in older versions of the Fabric OS might be disruptive to the fabric.

To customize the switch name

1. **SilkWorm 200E, 3250, 3850, 3900, 4100, 4900, 7500 switches, and SilkWorm 48000 directors**: Proceed to the next step.

   SilkWorm 24000 directors: Open a telnet window for each logical switch and enter the **switchName** command.

2. Connect to the switch and log in as admin.

3. **SilkWorm 200E, 3250, 3850, 3900, 4100, 4900, 7500 switches, and SilkWorm 48000 directors**: Proceed to the next step.

   SilkWorm 24000: If configured for one domain (the default) proceed to the next step. If configured with two domains, proceed as for the SilkWorm 24000.

4. Enter the **switchName** command at the command line, using the following syntax:

   ```
   switchname "newname"
   ```

   where *newname* is the new name for the switch.

5. Record the new switch name for future reference.

6. **SilkWorm 24000 configured with two domains**: Disconnect from the session and repeat the procedure for the second logical switch.

   ```
   switch:admin> switchname "switch62"
   Committing configuration...
   Done.
   switch62:admin>
   ```

Customizing the Chassis Name

Beginning with Fabric OS v4.4.0, it is recommended that you customize the chassis name for each switch. Some system logs identify switches by chassis names, so if you assign meaningful chassis names in addition to meaningful switch names, logs will be more useful.

To change the chassis name

1. Connect to the switch and log in as admin.
2. Enter the **chassisName** command at the command line, using the following syntax:

   ```
   switch:admin> chassisname "newname"
   ```

   where *newname* is the new name for the chassis.

   Chassis names can be from 1 to 15 characters long, must begin with a letter, and can contain letters, numbers, or the underscore character. It is not necessary to use the quotation marks.

3. Record the new chassis name for future reference.

Working With Domain IDs

Although domain IDs are assigned dynamically when a switch is enabled, you can reset them manually so that you can control the ID number or to resolve a domain ID conflict when you merge fabrics.

If a switch already has a domain ID when it is enabled, and that domain ID conflicts with a switch already in the fabric, the conflict is automatically resolved. The process can take several seconds, during which time traffic is delayed.

The default domain ID for SilkWorm switches is 1.

The default domain ID applies to the logical switches in SilkWorm 24000 directors that are configured for two domains. To prevent domain conflict, you can either disable one of the switches until the other is connected to the fabric, then reenable the switches so that unique domain IDs are automatically assigned, or you can use the procedure "To set the domain ID" on page 2-19 to make the domain IDs unique before connecting the logical switches to the fabric.

Caution
On switches running Fabric OS v4.0.0 and later, do not use domain ID 0, which is reserved for another purpose. The use of this domain ID can cause the switch to reboot continuously.

Avoid changing the domain ID on the FCS in secure mode. To minimize down time, change the domain IDs on the other switches in the secure fabric.

To display domain IDs

1. Connect to a switch and log in as admin.
2. Enter the **fabricShow** command.

   Fabric information is displayed, including the domain ID (D_ID):

```
switch:admin> fabricshow
Switch ID   Worldwide Name           Enet IP Addr    FC IP Addr      Name
-------------------------------------------------------------------------
  1: fffc01 10:00:00:60:69:e4:00:3c  10.32.220.80    0.0.0.0         "ras080"
  2: fffc02 10:00:00:60:69:e0:01:46  10.32.220.1     0.0.0.0         "ras001"
  3: fffc03 10:00:00:60:69:e0:01:47  10.32.220.2     0.0.0.0         "ras002"
  5: fffc05 10:00:00:05:1e:34:01:bd  10.32.220.5     0.0.0.0         "ras005"
  6: fffc06 10:00:00:05:1e:34:02:3e  10.32.220.6     0.0.0.0         "ras006"
  7: fffc07 10:00:00:60:69:34:02:0c  10.32.220.7     0.0.0.0         "ras007"
 10: fffc0a 10:00:00:60:69:80:04:46  10.32.220.10    10.32.219.0     "ras010"
 11: fffc0b 10:00:00:60:69:80:04:47  10.32.220.11    10.32.219.1     "ras011"
 15: fffc0f 10:00:00:60:69:80:47:74  10.32.220.15    0.0.0.0         "ras015"
 16: fffc10 10:00:00:60:69:80:47:75  10.32.220.16    0.0.0.0         "ras016"
 19: fffc13 10:00:00:05:1e:34:00:ad  10.32.220.19    0.0.0.0         "ras019"
 20: fffc14 10:00:00:05:1e:34:00:63  10.32.220.20    0.0.0.0        >"ras020"
 30: fffc1e 10:00:00:60:69:90:02:21  10.32.220.30    0.0.0.0         "ras030"
 31: fffc1f 10:00:00:60:69:90:02:60  10.32.220.31    0.0.0.0         "ras031"
 32: fffc20 10:00:00:60:69:90:02:68  10.32.220.32    0.0.0.0         "ras032"
 33: fffc21 10:00:00:60:69:90:03:20  10.32.220.33    0.0.0.0         "ras033"
 34: fffc22 10:00:00:60:69:90:03:01  10.32.220.34    0.0.0.0         "ras034"
 40: fffc28 10:00:00:60:69:50:06:7f  10.32.220.40    0.0.0.0         "ras040"
 45: fffc2d 10:00:00:05:1e:34:c5:17  10.32.220.45    0.0.0.0         "ras045"
 50: fffc32 10:00:00:60:69:c0:06:64  10.32.220.50    0.0.0.0         "ras050"
 51: fffc33 10:00:00:60:69:c0:1e:43  10.32.220.51    0.0.0.0         "ras051"
 60: fffc3c 10:00:00:60:69:12:34:44  10.32.220.60    0.0.0.0         "ras060"
 62: fffc3e 10:00:00:60:69:12:32:76  10.32.220.62    0.0.0.0         "ras062"
 63: fffc3f 10:00:00:60:69:12:45:6e  10.32.220.63    0.0.0.0         "ras063"
 64: fffc40 10:00:00:60:69:12:1d:51  10.32.220.64    0.0.0.0         "ras064"

The Fabric has 25 switches

switch:admin>
```

The fields in the **fabricShow** display are:

Switch ID—The switch Domain_ID and embedded port D_ID

Worldwide Name—The switch WWN

Enet IP Addr—The switch Ethernet IP address

FC IP Addr—The switch FC IP address

Name—The switch symbolic name. An arrow (>) indicates the principal switch.

To set the domain ID

1. Connect to the switch and log in as admin.
2. Enter the **switchDisable** command to disable the switch.
3. Enter the **configure** command.
4. Enter **y** after the "Fabric Parameters" prompt:

```
Fabric parameters (yes, y, no, n): [no] y
```

5. Enter a unique domain ID at the "Domain" prompt. Use a domain ID value from 1 through 239 for normal operating mode (FCSW compatible):

```
Domain: (1..239) [1] 3
```

6. Respond to the remaining prompts, or press **Ctrl-d** to accept the other settings and exit.
7. Enter the **switchEnable** command to reenable the switch.

Activating Ports on Demand

The SilkWorm 4900 model can be purchased with 32, 48, or 64 licensed ports. As your needs increase, you can activate unlicensed ports (up to a maximum of 64 ports) by purchasing and installing the Brocade Ports on Demand optionally licensed product.

The SilkWorm 4100 model can be purchased with 16, 24, or 32 licensed ports. As your needs increase, you can activate unlicensed ports (up to the maximum of 32 ports) by purchasing and installing the Brocade Ports on Demand optionally licensed product.

The SilkWorm 200E model can be purchased with 8 ports and no E_Port, 8 ports with full fabric access (SilkWorm 210E), and 16 ports with full fabric access (SilkWorm 240E). If you purchase the SilkWorm 200E with 8 ports enabled, you can activate unlicensed ports in 4-port increments up to 16 ports by purchasing and installing the Brocade Ports on Demand optional licensed product.You can also purchase an E_Port license upgrade if your switch does not yet support full fabric access.

Ports on Demand is ready to be unlocked in the switch firmware. Its license key might be part of the licensed Paper Pack supplied with switch software, or you can purchase the license key separately from your switch vendor. You might need to generate a license key from a transaction key supplied with your purchase. If so, launch an Internet browser and go to the Brocade Web site at *http://www.brocade.com*. Click **Products**> **Software Products**> **Software License Keys** and follow the instructions to generate the key.

By default, ports 0 through 15 are activated on the SilkWorm 4100. Each Ports on Demand license activates the next group of eight ports, in numerical order. Before installing a license key, you must insert transceivers in the ports to be activated. Remember to insert the transceivers in the lowest group of inactive port numbers first. For example, if only 16 ports are currently active and you are installing one Ports on Demand license key, make sure to insert the transceivers in ports 16 through 23. If you later install a second license key, insert the transceivers in ports 24 through 31. For details on inserting transceivers, see the *SilkWorm 4100 Hardware Reference Manual*.

After you install a license key, you must enable the ports to complete their activation. You can do so without disrupting switch operation by using the **portEnable** command on each port. Alternatively, you can disable and reenable the switch to activate ports.

To activate Ports on Demand

1. Connect to the switch and log in as admin.
2. *Optional*: To verify the current states of the ports, use the **portShow** command.

   In the **portShow** output, the Licensed field indicates whether the port is licensed or not.
3. Install the Brocade Ports on Demand license.

   For instructions, see "Maintaining Licensed Features" on page 2-14.
4. Use the **portEnable** command to enable the ports.
5. *Optional*: Use the **portShow** command to check the newly activated ports.

If you remove a Ports on Demand license, the licensed ports will become disabled after the next platform reboot or the next port deactivation.

Disabling and Enabling Switches

By default, the switch is enabled after power is applied and diagnostics and switch initialization routines have finished. You can disable and reenable it as necessary.

To disable a switch

1. Connect to the switch and log in as admin.
2. Enter the **switchDisable** command at the command line.

   All Fibre Channel ports on the switch are taken offline. If the switch was part of a fabric, the fabric is reconfigured.

To enable a switch

1. Connect to the switch and log in as admin.
2. Enter the **switchEnable** command at the command line.

   All Fibre Channel ports that passed the POST test are enabled. If the switch has interswitch links to a fabric, it joins the fabric.

Disabling and Enabling Ports

All licensed ports are enabled by default. You can disable and reenable them as necessary. Ports that you activate with Ports on Demand must be enabled explicitly, as described in "Activating Ports on Demand" on page 2-19.

Caution
The fabric will be reconfigured if the port you are enabling or disabling is connected to another switch.

To disable a port

1. Connect to the switch and log in as admin.
2. **SilkWorm 200E, 3250, 3850, 3900, 4100, 4900, and 7500 switches**: Enter the following command:

```
switch:admin> portdisable portnumber
```

where *portnumber* is the port number of the port you want to disable.

SilkWorm 24000 and 48000 directors: Enter the following command:

```
switch:admin> portdisable slotnumber/portnumber
```

where *slotnumber* and *portnumber* are the slot and port numbers of the port you want to disable.

To enable a port

1. Connect to the switch and log in as admin.
2. **SilkWorm 200E, 3250, 3850, 3900, 4100, 4900, and 7500 switches**: Enter the following command:

```
switch:admin> portenable portnumber
```

where *portnumber* is the port number of the port you want to enable.

SilkWorm 24000 and 48000 directors: Enter the following command:

```
switch:admin> portenable slotnumber/portnumber
```

where *slotnumber* and *portnumber* are the slot and port numbers of the port you want to enable. (Slots are numbered 1 through 4 and 7 through 10, counting from left to right.)

If the port is connected to another switch, the fabric might be reconfigured. If the port is connected to one or more devices, these devices become available to the fabric.

If you change port configurations during a switch failover, the ports might become disabled. To bring the ports online, reissue the **portEnable** command after the failover is complete.

Making Basic Connections

You can make basic connections to devices and to other switches.

Before connecting a switch to a fabric that contains switches running different firmware versions, you must first set the same PID format on all switches. The presence of different PID formats in a fabric causes fabric segmentation.

- For information on PID formats and related procedures, see "Selecting a PID format" on page A-4.
- For information on configuring the routing of connections, see "Routing Traffic" on page 10-1.
- For information on configuring extended interswitch connections, see "Administering Extended Fabrics" on page 18-1.

Connecting to Devices

To minimize port logins, power off all devices before connecting them to the switch. For devices that cannot be powered off, first use the **portDisable** command to disable the port on the switch, and then connect the device. When powering the devices back on, wait for each device to complete the fabric login before powering on the next one.

Connecting to Other Switches

See the hardware user's guide of your specific switch for interswitch link (ISL) connection and cable management information. The standard (default) ISL mode is L0, which you can configure with the **portCfgLongDistance** command. ISL Mode L0 is a static mode, with the following maximum ISL distances:

- 10 km at 1 Gbit/second
- 5 km at 2 Gbit/second
- 2.5 km at 4 Gbit/second

ISL mode L0 is available on all Fabric OS releases. When you upgrade from Fabric OS v4.0.0 to Fabric 4.1.0 or later, all extended ISL ports are set automatically to L0 mode.

For information on extended ISL modes, which enable longer distance interswitch links, see Chapter 18, "Administering Extended Fabrics".

Linking Through a Gateway

A gateway merges SANs into a single fabric by establishing point-to-point E_Port connectivity between two Fibre Channel switches that are separated by a network with a protocol such as IP or SONET.

Except for link initialization, gateways are transparent to switches; the gateway simply provides E_Port connectivity from one switch to another.

By default, switch ports initialize links using the Exchange Link Parameters (ELP) mode 1. However, gateways expect initialization with ELP mode 2, also referred to as ISL R_RDY mode. Therefore, to enable two switches to link through a gateway, the ports on both switches must be set for ELP mode 2.

Any number of E_Ports in a fabric can be configured for gateway links, provided the following rules are followed:

- All switches in the fabric must be upgraded to Fabric OS v3.1.0 (or later) or v4.1.0 (or later).
- To prevent fabric segmentation, make sure that all switches in the fabric are using the core PID format, as described in "To configure a link through a gateway" next.
- When determining switch count maximums, include the switches connected to both sides of the gateway.
- Extended links (those created using the Extended Fabrics licensed feature) and the security features in Secure Fabric OS are not supported through gateway links.

To configure a link through a gateway

1. If you are not sure that the PID format is consistent across the entire fabric, enter the **configShow** command on all switches to check the PID setting. If necessary, change the PID format on any nonconforming switches, as described in "Configuring the PID Format" on page A-1.
2. Connect to the switch on one end of the gateway and log in as admin.
3. Enter the **portCfgIslMode** command (see the *Fabric OS Command Reference* for more information about the **portCfgIslMode** command.):

   SilkWorm 200E, 3250, 3850, 3900, 4100, 4900, and 7500 switches: portCfgIslMode *port mode*

   Specify a *port* number. Valid values for port number vary depending on the switch type. The mode operand is required: specify 1 to enable ISL R_RDY mode (gateway link) or specify 0 to disable it.

   SilkWorm 24000 and 48000 directors: portCfgIslMode *slot*/*port*, *mode*

   Specify a *slot*/*port* number pair. Valid values for slot and port number vary depending on the switch type. The mode operand is required: specify 1 to enable ISL R_RDY mode (gateway link) or specify 0 to disable it.

   In the following example, slot 2, port 3 is enabled for a gateway link:

   ```
   switch:admin> portcfgislmode 2/3, 1
   Committing configuration...done.
   ISL R_RDY Mode is enabled for port 3. Please make sure the PID
   formats are consistent across the entire fabric.
   switch:admin>
   ```

4. Repeat the steps for any additional ports that will be connected to the gateway.
5. Repeat the procedure on the switch at the other end of the gateway.

Checking Status

You can check the status of switch operation, high availability features, and fabric connectivity.

To check switch operation

1. Connect to the switch and log in as admin.
2. Enter the **switchShow** command at the command line. This command displays a switch summary and a port summary.
3. Check that the switch and ports are online.
4. Use the **switchStatusShow** command to further check the status of the switch.

High Availability (HA) features provide maximum reliability and nondisruptive replacement of key hardware and software modules. To verify these features, connect to the switch as admin and use any of the following commands:

To verify High Availability features

1. Enter the **chassisShow** command to verify the field replaceable units (FRUs).

   SilkWorm directors:

   - **haShow** verifies that HA is enabled, that the heartbeat is up, and that the HA state is synchronized between the active and standby CP blades.
   - **slotShow** inventories and displays the current status of each slot in the system.

To verify fabric connectivity

1. Connect to the switch and log in as admin.
2. Enter the **fabricShow** command at the command line. This command displays a summary of all the switches in the fabric.

```
switch:admin> fabricshow
Switch ID   Worldwide Name           Enet IP Addr    FC IP Addr      Name
-------------------------------------------------------------------------
  1: fffc01 10:00:00:60:69:80:04:5a  192.168.186.61  192.168.68.193  "switch61"
  3: fffc03 10:00:00:60:69:10:9c:29  192.168.186.175 0.0.0.0         "switch175"
  4: fffc04 10:00:00:60:69:12:14:b7  192.168.174.70  0.0.0.0         "switch70"
  5: fffc05 10:00:00:60:69:45:68:04  192.168.144.121 0.0.0.0         "switch121"
  6: fffc06 10:00:00:60:69:00:54:ea  192.168.174.79  192.168.68.197  "switch79"
  7: fffc07 10:00:00:60:69:80:04:5b  192.168.186.62  192.168.68.194  "switch62"
  8: fffc08 10:00:00:60:69:04:11:22  192.168.186.195 0.0.0.0         "switch195"
  9: fffc09 10:00:00:60:69:10:92:04  192.168.189.197 192.168.68.198  "switch197"
 10: fffc0a 10:00:00:60:69:50:05:47  192.168.189.181 192.168.68.181  "switch181"
 11: fffc0b 10:00:00:60:69:00:54:e9  192.168.174.78  192.168.68.196  "switch78"
 15: fffc0f 10:00:00:60:69:30:1e:16  192.168.174.73  0.0.0.0         "switch73"
 33: fffc21 10:00:00:60:69:90:02:5e  192.168.144.120 0.0.0.0         "switch120"
 44: fffc2c 10:00:00:60:69:c0:06:8d  192.168.144.121 0.0.0.0         "switch121"
 97: fffc61 10:00:00:60:69:90:02:ed  192.168.144.123 0.0.0.0         "switch123"
 98: fffc62 10:00:00:60:69:90:03:32  192.168.144.122 0.0.0.0         "switch122"

The Fabric has 15 switches

switch:admin>
```

To verify device connectivity

1. Connect to the switch and log in as admin.
2. *Optional*: Enter the **switchShow** command to verify that devices, hosts, and storage are connected.
3. *Optional*: Enter the **nsShow** command to verify that devices, hosts, and storage have successfully registered with the Name Server.
4. Enter the **nsAllShow** command at the command line. This command displays 24-bit Fibre Channel addresses of all devices in the fabric.

```
switch:admin> nsallshow
{
  010e00 012fe8 012fef 030500 030b04 030b08 030b17 030b18
  030b1e 030b1f 040000 050000 050200 050700 050800 050de8
  050def 051700 061c00 071a00 073c00 090d00 0a0200 0a07ca
  0a07cb 0a07cc 0a07cd 0a07ce 0a07d1 0a07d2 0a07d3 0a07d4
  0a07d5 0a07d6 0a07d9 0a07da 0a07dc 0a07e0 0a07e1 0a0f01
  0a0f02 0a0f0f 0a0f10 0a0f1b 0a0f1d 0b2700 0b2e00 0b2fe8
  0b2fef 0f0000 0f0226 0f0233 0f02e4 0f02e8 0f02ef 210e00
  211700 211fe8 211fef 2c0000 2c0300 611000 6114e8 6114ef
  611600 620800 621026 621036 6210e4 6210e8 6210ef 621400
  621500 621700 621a00
75 Nx_Ports in the Fabric }
switch:admin>
```

The number of devices listed should reflect the number of devices that are connected.

Tracking and Controlling Switch Changes

The track changes feature allows you to keep record of specific changes that might not be considered switch events, but might provide useful information. The output from the track changes feature is dumped to the system messages log for the switch. Use the **errDump** or **errShow** command to view the log.

Items in the log created from the Track changes feature are labeled *Track*.

Trackable changes are:

- Successful login
- Unsuccessful login
- Logout
- Configuration file change from task
- Track changes on
- Track changes off

An SNMP-TRAP mode can also be enabled (see the **trackChangesHelp** command in the *Fabric OS Command Reference*).

For troubleshooting information on the track changes feature, see “Inaccurate Information in the System Message Log” on page 15-25.

To enable the track changes feature

1. Connect to the switch and log in as admin.
2. Enter this command to enable the track changes feature: **trackChangesSet 1**.

   A message displays, verifying that the track changes feature is on:

```
switch:admin> trackchangesset 1
Committing configuration...done.
switch:admin>
```

   The output from the track changes feature is dumped to the system message log for the switch.
3. Use the **errDump** or **errShow** command to view the log.

   Items in the system message log created from the track changes feature are labeled TRCK:

```
2004/08/24-08:45:43, [TRCK-1001], 212,, INFO, ras007, Successful login by user
admin.
```

To display the status of the track changes feature

1. Connect to the switch and log in as admin.
2. Enter the **trackChangesShow** command.

   The status of the track changes feature is displayed as either on or off. The display includes whether or not the track changes feature is configured to send SNMP traps:

```
switch:admin> trackchangesshow
Track changes status: ON
Track changes generate SNMP-TRAP: NO
switch:admin>
```

To view the switch status policy threshold values

1. Connect to the switch and log in as admin.
2. Enter the **switchStatusPolicyShow** command at the command line.

   Whenever there is a switch change, an error message is logged and an SNMP connUnitStatusChange trap is sent.

   SilkWorm 200E, 3250, 3850, 3900, 4100, 4900, and 7500 switches: The output is similar to the following:

```
switch:admin> switchstatuspolicyshow
The current overall switch status policy parameters:
                  Down    Marginal
 ---------------------------------
     PowerSupplies  2         1
      Temperatures  2         1
              Fans  2         1
             Flash  0         1
     MarginalPorts  5         2
       FaultyPorts  2         1
       MissingSFPs  2         1
switch:admin>
```

SilkWorm 24000 and 48000: The output is similar to the following:

```
switch:admin> switchstatuspolicyshow
The current overall switch status policy parameters:
                   Down    Marginal
 ---------------------------------
     PowerSupplies  3         0
      Temperatures  2         1
              Fans  2         1
               WWN  0         1
                CP  0         1
             Blade  0         1
             Flash  0         1
     MarginalPorts  2         1
       FaultyPorts  2         1
       MissingSFPs  0         0
switch:admin>
```

The policy parameter determines the number of failed or inoperable units for each contributor that will trigger a status change in the switch.

Each parameter can be adjusted so that a specific threshold must be reached before that parameter changes the overall status of a switch to MARGINAL or DOWN. For example, if the FaultyPorts DOWN parameter is set to 3, the status of the switch will change if 3 ports fail. Only one policy parameter needs to pass the MARGINAL or DOWN threshold to change the overall status of the switch.

For more information about setting policy parameters, see the *Fabric Watch Administrator's Guide*.

To set the switch status policy threshold values

1. Connect to the switch and log in as admin.
2. Enter the **switchStatusPolicySet** command at the command line.

   The current switch status policy parameter values are displayed. You are prompted to enter values for each DOWN and MARGINAL threshold parameter.

Note

By setting the DOWN and MARGINAL value for a parameter to *0,0* that parameter is no longer used in setting the overall status for the switch.

3. Verify the threshold settings you have configured for each parameter.

   Enter the **switchStatusPolicyShow** command to view your current switch status policy configuration.

   SilkWorm 3250, 3850, 3900, 4100, 4900, and 7500 switches: The following example shows the command as executed on a SilkWorm 3900 switch. The output is similar on SilkWorm 3250, 3850, 4100, 4900, and 7500 models:

```
switch:admin> switchstatuspolicyset
To change the overall switch status policy parameters

The current overall switch status policy parameters:
                   Down    Marginal
 ---------------------------------
     PowerSupplies    2           1
      Temperatures    2           1
              Fans    2           1
             Flash    0           1
     MarginalPorts    2           1
       FaultyPorts    2           1
       MissingSFPs    0           0

Note that the value, 0, for a parameter, means that it is
NOT used in the calculation.
** In addition, if the range of settable values in the prompt is (0..0),
** the policy parameter is NOT applicable to the switch.
** Simply hit the Return key.

Bad PowerSupplies contributing to DOWN status: (0..2) [2] 0
Bad PowerSupplies contributing to MARGINAL status: (0..2) [1] 0
Bad Temperatures contributing to DOWN status: (0..6) [2] 0
Bad Temperatures contributing to MARGINAL status: (0..6) [1] 0
Bad Fans contributing to DOWN status: (0..3) [2] 0
Bad Fans contributing to MARGINAL status: (0..3) [1] 0
Out of range Flash contributing to DOWN status: (0..1) [0] 0
Out of range Flash contributing to MARGINAL status: (0..1) [1] 0
MarginalPorts contributing to DOWN status: (0..32) [2] 0
MarginalPorts contributing to MARGINAL status: (0..32) [1] 0
FaultyPorts contributing to DOWN status: (0..32) [2] 0
FaultyPorts contributing to MARGINAL status: (0..32) [1] 0
MissingSFPs contributing to DOWN status: (0..32) [0] 0
MissingSFPs contributing to MARGINAL status: (0..32) [0] 0

Policy parameter set has been changed
rint12:admin>
```

   SilkWorm 24000 and 48000: Command output includes parameters related to CP blades.

Configuring the Audit Log

When managing SANs, you may wish to filter, or audit, certain classes of events to ensure that you can view and generate a paper trail, or “audit log,” for what is happening on a switch, particularly for security elated event changes. These events include login failures, zone configuration changes, firmware downloads, and other configuration changes—in other words—critical changes that have a serious effect on the operation and security of the switch.

Important information related to event classes is also tracked and made available. For example, you can track changes from an external source via user name, IP address, or type of management interface used to access the switch.

Auditable events are generated by the switch and streamed to an external host through a configured system message log daemon (syslog). You specify a filter on the output to select the event classes that are sent through the system message log. The filtered events are streamed chronologically and sent to the system message log on an external host in the specified audit message format. This ensures that they can be easily distinguished from other system message log events that occur in the network. Then, at some regular interval of your choosing, you can review the audit events to look for unexpected changes.

Before you configure audit event logging, familiarize yourself with the following audit event log behaviors and limitations:

- By default, *all event classes* are configured for audit; to create an audit event log *for specific events*, you must explicitly set a filter via the *class* operand and then enable it.
- Audited events are generated specific to a switch and have no negative impact on performance.
- All Secure Fabric OS event are audited.
- Events are not persistently stored on the switch but are streamed to a system message log.
- The audit log depends on the system message log facility and IP network to send messages from the switch to a remote host. Because the audit event log configuration has no control over these facilities, audit events can be lost if the system message log and IP network facilities fail.
- If too many events are generated by the switch, the system message log will become a bottleneck and audit events will be dropped by the Fabric OS.
- If the user name, IP address, or user interface is not transported an audit message is logged by adding the message None to each of the respective fields.
- For High Availability, the audit event logs exist independently on both active and standby CPs. The configuration changes that occur on the active CP are propagated to the standby CP and take effect.
- Audit log configuration is updated via a configuration download.

See the *Fabric OS Command Reference* for more information about the **auditCfg** command and command syntax.

Auditable Event Classes

You configure the audit log using the **auditCfg** command. Before configuring an audit log, you must select the event classes you want audited. When enabled, the audit log feature audits any RASLOG messages (system message log) previously tagged as AUDIT in Fabric OS v5.1.0, which includes:

- SEC-3001 through SEC-3017
- SEC-3024 through SEC-3029
- ZONE-3001 through ZONE-3012

The following table identifies auditable event classes and **auditCfg** operands used to enable auditing of a specific class.

Table 2-2 AuditCfg Event Class Operands

Operand	Event Class	Description
1	Zone	Audit zone event configuration changes, but not the actual values that were changed. For example, you a message might state, “Zone configuration has changed,” but the syslog does not display the actual values that were changed.
2	Security	Audit any user-initiated security event for all management interfaces. For events that have an impact on an entire fabric, an audit is generated only for the switch from which the event was initiated.
3	Configuration	Audit configuration downloads of existing SNMP configuration parameters. Configuration uploads are not audited.
4	Firmware	Audit firmware download start, firmware complete, and any other errors encountered during a firmware download.
5	Fabric	Audit administrative domain-related changes.

Note

Only the active CP can generate audit messages because event classes being audited occur only on the active CP. Audit messages cannot originate from other blades in a chassis.

Audit events have the following message format:

```
AUDIT, <Timestamp>, [<Event ID>], <Severity>, <Event Class>, <User ID>/<Role>/<IP
address>/<Interface>,<Admin Domain>/<Switch name>,<Reserved>,<Event-specific
information>
```

Switch names are logged for switch components and chassis names for chassis components. For example, a chassis name might be FWDL or RAS and a switch component name might be zone, name server, or SNMP.

Pushed messages contain the administration domain of the entity that generated the event. See the *Fabric OS Message Reference* for details on message formats. See Chapter 14, “Working With Diagnostic Features” for details on setting up the system error log daemon.

Audit logging assumes that your syslog is operational and running. Before configuring an audit log, you must perform the following steps to ensure that the host syslog is operational and running.

To verify host syslog prior to configuring the audit log

1. Set up an external host machine with a system message log daemon running to receive the audit events that will be generated.
2. On the switch where the audit configuration is enabled, enter the **syslogdipaddrAdd** command to add the IP address of the host machine so that it can receive the audit events.
3. Ensure the network is configured with a network connection between the switch and the remote host.
4. Check the host SYSLOG configuration. If all error levels are not configured, you may not see some of the audit messages.

To configure an audit log for specific event classes

1. Connect to the switch from which you wish to generate an audit log and log in as admin.
2. Enter the **auditCfg --class** command, which defines the specific event classes to be filtered.

```
switch:admin> auditcfg --class 2,4
Audit filter is configured.
```

   The **auditCfg** event class operands are identified in Table 2-2

3. Enter the **auditCfg --enable** command, which enables audit event logging based on the classes configured in Step 2.

```
switch:admin> auditcfg --enable
Audit filter is enabled.
```

   To disable an audit event configuration, enter the **auditCfg --disable** command.

4. Enter the **auditCfg --show** command to view the filter configuration and confirm that the correct event classes are being audited, and the correct filter state appears (enabled or disabled).

```
switch:admin> auditcfg --show
 Audit filter is enabled.
 2-SECURITY
 4-FIRMWARE
```

5. To verify the audit event log setup, make a change affecting an enabled event class, and confirm that the remote host machine receives the audit event messages.

   The following example shows the SYSLOG (system message log) output for audit logging.

```
Jun  2 08:33:04 [10.32.220.7.2.2] raslogd: AUDIT, 2006/06/02-15:25:53, [SULB-
1003], INFO, FIRMWARE, root/root/NONE/console/CLI, ad_0/ras007_chassis, ,
Firmwarecommit has started.
Jun 5 06:45:33 [10.32.220.70.2.2] raslogd: AUDIT, 2006/06/05-13:38:17, [CONF-
1010], INFO, CONFIGURATION, root/root/NONE/
console/CLI, ad_0/ras070, , configDownload failed
Jun 5 08:15:32 [10.32.248.73.2.2] raslogd: AUDIT, 2006/06/05-13:38:17, [SEC-
1000], WARNING, SECURITY, JaneDoe/root/192.168.132.19/
telnet, Domain A/DoeSwitch, , Incorrect password during login attempt.
```

Shutting Down Switches and Directors

To avoid corrupting your file system, it is recommended that you perform graceful shutdowns of SilkWorm switches and directors. For directors running Fabric OS versions *prior* to v5.1.0, the following procedure describe how to gracefully shut down a director:

To power off a SilkWorm director gracefully (prior to v5.1.0)

1. Verify which CP is the active CP, and log in to the active CP using a Serial Console connection.
2. On the standby CP, set the slider switch to the off position, or eject the standby CP from the chassis. This disables the standby CP.
3. Enter the **reboot** command from the active CP. This will gracefully take down the system.
4. When you see the “Press escape within 4 seconds to enter boot interface” message, press ESC to suspend the active CP.
5. Power off the chassis by flipping both AC power switches to “0” (LEDs inside AC power switches should turn off). To maintain the ground connection, leave both power cords connected to the chassis and to an electrical outlet.

For both switches and directors running Fabric OS v5.1.0 and later, it is recommended that you use the following graceful shutdown procedures.

To power off a SilkWorm switch gracefully (v5.1.0 and later)

1. Connect to the switch and log in as admin.
2. Enter the **sysShutdown** command.
3. At the prompt, enter **y**.

```
switch:admin> sysshutdown
This command will shutdown the operating systems on your switch.
You are required to power-cycle the switch in order to restore operation.
Are you sure you want to shutdown the switch [y/n]?y
```

4. Wait until the following message displays:

```
Broadcast message from root (ttyS0) Wed Jan 25 16:12:09 2006...

The system is going down for system halt NOW !!
INIT: Switching to runlevel: 0
INIT: Sending processes the TERM signal
Unmounting all filesystems.
The system is halted
flushing ide devices: hda
Power down.
```

5. Power off the switch.

To power off a SilkWorm director gracefully (v5.1.0 and later)

1. From the active CP in a dual CP platform, enter the **sysShutdown** command.

Note

When the **sysShutdown** command is issued on the active CP, the active CP, the standby CP, and any AP blades are all shut down.

2. At the prompt, enter **y**.
3. Wait until you see the following message:

```
2006/01/25-17:01:40, [FSSM-1003], 194,, WARNING, NANDU, HA State out of sync
HA is disabled

Broadcast message from root (ttyS0) Wed Jan 25 17:01:41 2006...

The system is going down for system halt NOW !!
INIT: Switching to runlevel: 0
INIT: Sending processes the TERM signal
Unmounting all filesystems.
The system is halted
flushing ide devices: hda
Power down.
```

4. Power off the switch.

High Availability of Daemon Processes

Fabric OS v5.2.0 supports automatic restart of non-critical daemons. Starting these non-critical daemons is automatic, you cannot configure the startup process. The following sequence of events occur when a non-critical daemon fails:

1. When a non-critical daemon fails or dies, a RASlog and AUDIT event message is logged.
2. The daemon is automatically started again.
3. If the re-start is successful, then another message is sent to RASlog and AUDIT reporting the successful restart status.
4. Should the restart fail, another message is sent to RASlog and no further attempts are made to re-start the daemon.

Schedule downtime and reboot the switch at your convenience. The following table lists the daemons that are considered non-critical and are automatically restarted on failure.

Table 2-3 List of Daemons That Are Automatically Restarted

Daemon	Description
Snmpd	Simple Network Management Protocol Daemon.
Webd	Webserver daemon used for WebTools (includes httpd as well).
Cald	Common Access Layer Daemon (used by Manageability Applications).
Raslogd	Remote Access Service Log Daemon logs error detection, reporting, handling, and presentation of data into a format consumable by management tools and the user.
Rpcd	Remote Procedure Call daemon used by the API (Fabric Access API and SMI-S).
Evmd	Event Monitor Daemon (Port and Switch SCNs, firmwareDownload, configDownload).
Arrd	Asynchronous Response Router (used to send management data to hosts when the switch is accessed via the APIs (FA API or SMI-S).
Trackd	Track Changes Daemon.

Managing User Accounts

Chapter 3

This chapter provides information and procedures on managing authentication and user accounts for the switch management channel.

This chapter covers the following topics:

- "Overview" on page 3-1
- "Configuring the Authentication Model" on page 3-6
- "Managing the Local Database User Accounts" on page 3-7
- "Configuring the Local User Database" on page 3-12
- "Distributing the Local User Database" on page 3-12
- "Managing Fabric OS Users on the RADIUS Server" on page 3-17
- "Setting the Boot PROM Password" on page 3-30
- "Recovering Forgotten Passwords" on page 3-35

Overview

Fabric OS provides two options for authenticating users—remote RADIUS services and/or the local switch user database. Both options allow users to be centrally managed using the following methods:

- **Local user database**: Manually synchronize the local user database using the distribute command to push a copy of the switch's local user database to all other Fabric OS v5.2.0 and higher switches in the fabric.
- **Remote RADIUS servers**: Users are managed in a remote RADIUS server, all switches in the fabric can be configured to authenticate against the centralized remote database.

Accessing the Management Channel

The total number of sessions on a switch may not exceed 32. Table 3-1 shows the number of simultaneous login sessions allowed for each role.

Table 3-1 Maximum Number of Simultaneous Sessions

Role Name	Maximum Sessions
User	4
Operator	4
SwitchAdmin	4
ZoneAdmin	4
FabricAdmin	4
BasicSwitchAdmin	4
Admin	2

Using Role-Based Access Control (RBAC)

Fabric OS v5.2.0 uses Role-Based Action Control (RBAC) to determine which commands a user can run. Assign one of the Fabric OS predefined roles to a user, as shown in Table 3-2.

Table 3-2 Fabric OS v5.2.0 Roles

Role name	Fabric OS Version	Duties	Description
Admin	All	All administration	All administrative commands.
SwitchAdmin	v5.0.0 and higher	Local switch administration	Most switch (local) commands, excludes security, user management, and zoning commands.
Operator	v5.2.0 and higher	General switch administration	Routine switch maintenance commands.
ZoneAdmin	v5.2.0 and higher	Zone administration	Zone management commands only.
FabricAdmin	v5.2.0 and higher	Fabric and switch administration	All switch and fabric commands, excludes user management and Administrative Domains commands.
BasicSwitchAdmin	v5.2.0 and higher	Restricted switch administration	Mostly monitoring with limited switch (local) commands.
User	All	Monitoring only	Nonadministrative use, such as monitoring system activity.

Role Permissions

Table 3-3 describes the types of permissions that are assigned to roles.

Table 3-3 Permission Types

Abbreviation	Definition	Description
O	Observe	The user can run commands using options that display information only, such as running **userConfig --show -a** to show all users on a switch.
M	Modify	The user can run commands using options that create, change, and delete objects on the system, such as running **userconfig --change *username* -r *rolename*** to change a user's role.
OM	Observe-Modify	The user can run commands using both observe and modify options; if a role has modify permissions, it almost always has observe.
N	None	The user is not allowed to run commands in that category.

The Table 3-2 shows the permission type for categories of commands that each role is assigned. The permissions apply to all commands within the specified category. For a complete list of commands and role permissions, see *Fabric OS Command Reference*.

Table 3-4 RBAC Permissions Matrix

Category	Role Permission						
	User	Operator	Switch Admin	Zone Admin	Fabric Admin	Basic Switch Admin	Admin
Admin Domains	N	N	N	N	N	N	OM
Admin Domains—Selection	OM	OM	OM	OM	OM	OM	OM
APM	O	O	OM	N	OM	O	OM
Audit	O	O	O	O	O	O	O
Authentication	N	N	N	N	N	N	OM
Blade	O	OM	OM	N	OM	O	OM
Chassis Configuration	O	OM	OM	N	OM	O	OM
Configuration Management	N	O	O	O	O	O	OM
Debug	N	N	N	N	N	N	N
Diagnostics	O	OM	OM	N	OM	O	OM
Ethernet Configuration	O	O	OM	N	OM	O	OM
Fabric	O	O	O	N	OM	O	OM
Fabric Distribution	N	N	N	N	OM	N	OM
Fabric Routing	O	O	O	O	OM	O	OM
Fabric Watch	O	OM	OM	N	OM	O	OM

Table 3-4 RBAC Permissions Matrix (Continued)

Category	Role Permission						
	User	Operator	Switch Admin	Zone Admin	Fabric Admin	Basic Switch Admin	Admin
FICON	O	OM	OM	N	OM	O	OM
Firmware Management	O	OM	OM	O	OM	O	OM
FRU Management	O	OM	OM	N	OM	O	OM
HA (High Availability)	O	O	OM	N	OM	O	OM
iSCSI	O	O	O	O	OM	O	OM
Switch Management—IP Configuration	O	OM	OM	N	OM	O	OM
Local User Environment	OM	OM	OM	OM	OM	OM	OM
Logging	O	OM	OM	O	OM	O	OM
License	O	OM	OM	O	OM	O	OM
Management Access Configuration	O	OM	OM	N	OM	O	OM
Management Server	O	OM	OM	O	OM	O	OM
Name Server	O	O	OM	O	OM	O	OM
Nx_Port Management	O	O	OM	O	OM	O	OM
Physical Computer System	O	O	O	N	O	O	O
PKI	O	O	O	N	O	O	OM
Port Mirroring	N	N	N	N	N	N	OM
RADIUS	N	N	N	N	N	N	OM
Routing—Basic	O	OM	OM	O	OM	O	OM
Routing—Advanced	O	O	O	N	OM	O	OM
Security	O	N	O	N	OM	O	OM
Session Management	O	OM	OM	N	OM	OM	OM
SNMP	O	O	OM	N	OM	O	OM
Statistics	O	OM	OM	N	OM	O	OM
Statistics—Device	O	OM	OM	N	OM	O	OM
Statistics—Port	O	OM	OM	N	OM	O	OM
Switch Management	O	OM	OM	O	OM	O	OM
Switch Configuration	O	OM	OM	N	OM	O	OM

Table 3-4 RBAC Permissions Matrix (Continued)

Category	Role Permission						
	User	Operator	Switch Admin	Zone Admin	Fabric Admin	Basic Switch Admin	Admin
Switch Port Configuration	O	OM	OM	N	OM	OM	OM
Switch Port Management	O	OM	OM	O	OM	OM	OM
Topology	O	O	O	N	OM	O	OM
User Management	N	N	N	N	N	N	OM
WWN Card	O	OM	OM	N	OM	N	OM
Zoning	O	O	O	OM	OM	O	OM

Configuring the Authentication Model

This section explains how to configure authentication of the switch management channel connections. Fabric OS v5.2.0 and higher supports use of both the local user database and RADIUS service at the same time. Use the **aaaConfig** command to set the authentication model for Fabric OS switch management channel connection authentication model as shown in Table 3-5.

Note

Set the authentication model on each switch.

Table 3-5 Authentication Configuration Options

aaaConfig Option	Description	Equivalent Setting in Fabric OS v5.1.0 and older	
		--radius	--switchdb[a]
--localonly	Default setting. Authenticates management connections against the local database only. If the password does not match or the user is not defined, the login fails.	Off	On
--radiusonly[b]	Authenticates management connections against the RADIUS database(s) only. If the RADIUS service is not available or the credentials do not match, the log in fails.	On	Off
--radiuslocal	Authenticates management connections against any RADIUS databases first. If RADIUS fails *for any reason*, authenticates against the local user database.	not supported	not supported
--radiuslocalbackup	Authenticates management connections against any RADIUS databases. If RADIUS fails because the service is not available, authenticates against the local user database.	On	On

a. Fabric OS v5.1.0 and earlier aaaConfig --switchdb <on | off> setting.

b. The console login will never be set to **--radiusonly** mode for login recovery purposes. When **--radiusonly** mode is turned on, console login uses the **--radiuslocalbackup** mode.

To set the switch authentication model

1. Connect to the switch and log in.
2. Enter this command:

```
switch:admin> aaaConfig [--localonly | --radiusonly | --radiuslocal |
--radiuslocalbackup]
```

Managing the Local Database User Accounts

User add, change, and delete operations are subject to the *subset* rule: an admin with ADlist 0-10 cannot perform operations on an *admin*, *user*, or *any* role with an ADlist 11-25. The user account being changed must have an ADlist that is a subset of the account that is making the change.

About the Default Accounts

Fabric OS provides the following predefined accounts in the switch-local user database. Change the password for all defaults during the initial installation and configuration.

Table 3-6 Default Local User Accounts

Account Name	Role	Admin Domain	Description
user	User	AD0 home: 0	Most commands have *observe-only* permission.
admin	Admin	AD0-255 home: 0	Most commands have *observe-modify* permission.
root	Root	AD0-255 home: 0	Reserved.
factory	Factory	AD0-255 home: 0	Reserved.

Defining Local User Accounts

In addition to the default administrative and user accounts, Fabric OS supports up to 256 user-defined accounts in each logical switch (domain). These accounts expand your ability to track account access and audit administrative activities.

The following procedures can be performed on local user accounts. Administrators can act on other accounts only if that account has an Admin Domain list that is a subset of the administrator.

Note

When operating in secure mode, you must perform these operations on the primary FCS switch. The userConfig command with Admin Domain-related options is not valid in secure mode.

To display account information

1. Connect to the switch and log in.
2. Enter one of the “show” commands:
   - **userConfig --show -a** to show all account information for a logical switch
   - **userConfig --show -b** to show all backup account information for a logical switch
   - **userConfig --show username** to show account information for the specified account
   - **userConfig --showad -a admindomain_ID** to show all accounts permitted to select the specified admindomain_ID

To create an account

1. Connect to the switch and log in.
2. Enter the following command:

```
userConfig --add <username> -r <rolename> [-h <admindomain_ID>]
[-a <admindomain_ID list>] [-d <description>] [-x]
```

username	Specifies the account name, which must begin with an alphabetic character. The name can be from 8 to 40 characters long. It is case-sensitive and can contain alphabetic and numeric characters, the dot (.) and the underscore (_). It must be different than all other account names on the logical switch. The account name cannot be the same as a role name.
-r *rolename*	Specifies the role: either User, SwitchAdmin, ZoneAdmin, FabricAdmin, BasicSwitchAdmin, Operator, or Admin in nonsecure mode; in secure mode you can also use NonfcsAdmin.
-h *admindomain_ID*	*Optional:* Specifies the home Administrative Domain; if no Administrative Domain is specified, then lowest numbered Administrative Domain in the list is assigned.
-a *admindomain_ID*	*Optional:* Specifies which Administrative Domain(s) the user may access; if no Administrative Domains are listed the user is automatically assigned to AD0. Use comma-separated lists and/or ranges, for example -a 0,9,10-15,244.
-d *description*	*Optional:* Adds a description to the account. The description field can be up to 40 printable ASCII characters long. The following characters are not allowed: asterisk (*), single quote mark (‘), quotation mark (“), exclamation point (!), semicolon (;), and colon (:).
-x	Specifies an expired password that must be changed the first time the user logs in.

3. In response to the prompt, enter a password for the account.

   The password is not displayed when you enter it on the command line.

To delete an account

1. Connect to the switch and log in.
2. Enter the following command:

```
userConfig --delete username
```

where *username* specifies the account name.

Note

You cannot delete the default accounts. An account cannot delete itself. All active CLI sessions for the deleted account are logged out.

3. At the prompt for confirmation, enter **y**.

To change account parameters

When changing account parameters, if you change the ADlist for the user account, all of the currently active sessions for that account will be logged out.

1. Connect to the switch and log in.
2. Enter the following command:

```
userconfig --change username [-r rolename] [-h admindomain_ID]
[-a admindomain_ID_ list] [-d description] [-e yes | no] -u -x
```

username	Changes the account attribute for username. The account must already exist.
-r *rolename*	Optionally, changes the role to one of the names listed in Table 3-2 on page 3-2. In secure mode, role can also be changed to nonfcsadmin role. An account cannot change its own role. Account with Admin role can change the role names of all user-defined accounts except those with Admin roles.
-h *admindomain_ID*	Changes the home Administrative Domain.
-a *admindomain_ID_list*	Adds the new list Administrative Domain to the existing list.
-d *description*	Optionally, the account description. The description field can be up to 40 printable ASCII characters long. The following characters are not allowed: asterisk (‘), quotation mark (“), exclamation point (!), semicolon (;), and colon (:).
-e	Optionally, enter **yes** to enable the account or enter **no** to disable it. If you disable an account, all active CLI sessions for that account are logged out. You can enable or disable user-defined or default accounts.
-u	Unlocks the user account.
-x	Specifies an expired password that must be changed the first time the user logs in.

To add an Administrative Domain to the account

1. Connect to the switch and log in.
2. Enter the command

   userConfig --addad <username> [-h <admindomain_ID>] [-a <admindomain_ID_list>]

   where <username> is the Administrative Domain attribute for the username: the account must already exist, <admindomain_ID> is the home Administrative Domain and <admindomain_ID_list> adds the new list Administrative Domain to the existing list.
3. Log into the switch again to access the newly-added Admin Domain.

To remove an Administrative Domain from the Account

When removing an Admin Domain from an account, all of the currently active sessions for that account will be logged out.

1. Connect to the switch and log in.
2. Enter the command

   userConfig --deletead <username> [-h <admindomain_ID>] [-a <admindomain_ID_list>]

   where <username> is the new Admin Domain attribute for username: the account must already exist, admindomain_ID is the home Admin Domain, and admindomain_ID_list is the Admin Domain list to be removed from the existing list. If the **–h** argument is not specified, the home Admin Domain will either remain as it was or will be the lowest Admin Domain ID in the remaining list.

Recovering accounts

The following conditions apply to recovering user accounts:

- The attributes in the backup database replace the attributes in the current account database.
- An event is stored in the system message log, indicating that accounts have been recovered.

In secure mode, fabric distribution overwrites the local switch account database. The original local switch database is saved as a backup and can be recovered. The local switch database is not backed up in nonsecure mode.

To recover an account

1. Connect to the switch and log in.
2. If a backup account exists, enter the following command.

   userConfig --recover

   The AD list for a user account is not recovered; recovered accounts are given access only to AD0, regardless of previous AD assignments

Changing Local Account Passwords

The following rules apply to changing passwords:

- A user can change their own password.
- Only users with Admin roles can change the password for other accounts. When changing an Admin account password, you must provide the current password.
- An admin with ADlist 0-10 cannot change the password on an *admin*, *user*, or *any* role with an ADlist 11-25. The user account being changed must have an ADlist that is a subset of the account that it making the change.
- A new password must have at least one character different from the old password.
- You cannot change passwords using SNMP.
- Password prompting is disabled when security mode is enabled.

Note

Starting with Fabric OS v4.4.0, accounts with the Admin role can use Web Tools to change passwords.
Starting with Fabric OS v3.2.0, you cannot change default account names.
Starting with Fabric OS v5.1.0 password policies apply.

For information on password behavior when you upgrade (or downgrade) firmware, see "Upgrading and Downgrading Firmware" on page 8-2.

To change the password for the current login account

1. Connect to the switch and log in.
2. Enter the following command:

```
passwd
```

   Enter the requested information at the prompts.

To change the password for a different account

1. Connect to the switch and log in.
2. Enter the following command:

```
passwd name
```

   where *name* is the name of the account.

   Enter the requested information at the prompts.

Configuring the Local User Database

This section covers the following topics:

- "Distributing the Local User Database" on page 3-12
- "Protecting the Local User Database from Distributions" on page 3-13
- "Configuring Password Policies" on page 3-13

Distributing the Local User Database

This section explains how to distribute the switch-local user database and passwords to other switches in the fabric. The distributed database replaces (overwrites) the database on the target switch. The 'Locked' status of a user account is not distributed as part of local user database distribution.

When distributing the user database, the database may be rejected for one of the following reasons:

- One of the target switches does not have Fabric OS v5.2.0 or higher.
- One of the target switches user database is protected.

Caution
Distribute the user database and password policies only to Fabric OS v5.2.0 or higher switches; the distribution command fails if any of the targets is an earlier version.

To distribute the local user database

When distributing the local user database, all user-defined accounts residing in the receiving switches will be logged out of any active sessions.

1. Connect to the switch.
2. Enter the following command:

```
distribute -p PWD -d <switch_list>
```

where <switch_list> is a semicolon-separated list of switch Domain IDs, switch names, or switch WWN addresses. You can also specify **-d *** to sent the local user database only to Fabric OS v5.2.0 or higher switches in the fabric.

Protecting the Local User Database from Distributions

Fabric OS v5.2.0 and higher allows you to distribute the user database and passwords to other switches in the fabric. When the switch accepts a distributed user database, it replaces the local user database with the user database it receives.

By default, Fabric OS v5.2.0 switches accept the user databases and passwords distributed from other switches. This section explains how to protect the local user database from being overwritten.

To accept the user database

1. Connect to the switch.
2. Enter the following command:

```
fddCfg --localaccept PWD
```

   where PWD is one of the three supported database policies. Supported policy databases are SCC, DCC, PWD.

To reject distributed user databases

1. Connect to the switch.
2. Enter the following command:

```
fddCfg --localreject PWD
```

Configuring Password Policies

The password policies described in this section apply to the switch-local user database only. Configured password policies (and all user account attribute and password state information) are synchronized across CPs and remain unchanged after an HA failover. Password policies can also be manually distributed across the fabric (see “Distributing the Local User Database” on page 3-12).

All password policies are enforced during logins to the standby CP. However, you might observe that the password enforcement behavior on the standby CP is inconsistent with prior login activity because password state information from the active CP is automatically synchronized with the standby CP, thereby overwriting any password state information that was previously stored there. Also, password changes are not permitted on the standby CP.

Password authentication policies configured using the **passwdCfg** command are *not* enforced during initial prompts to change default passwords.

- Password strength
- Password history
- Password expiration

- Account lockout

Note

Secure mode supports only the default values of the password policies. If you attempt to enable secure mode after configuring changing any of the password policies you receive an error.

Setting the Password Strength Policy

The password strength policy is enforced across all user accounts, and enforces a set of format rules to which new passwords must adhere. The password strength policy is enforced only when a new password is defined. Use the following attributes to set the password strength policy:

- Lowercase

  Specifies the minimum number of lowercase alphabetic characters that must appear in the password. The default value is zero. The maximum value must be less than or equal to the MinLength value.

- Uppercase

  Specifies the minimum number of uppercase alphabetic characters that must appear in the password. The default value is zero. The maximum value must be less than or equal to the MinLength value.

- Digits

  Specifies the minimum number of numeric digits that must appear in the password. The default value is zero. The maximum value must be less than or equal to the MinLength value.

- Punctuation

  Specifies the minimum number of punctuation characters that must appear in the password. All printable, non-alphanumeric punctuation characters except colon ":" are allowed. The colon character not allowed because it is incompatible Brocade WebTools. The default value is zero. The maximum value must be less than or equal to the MinLength value.

- MinLength

  Specifies the minimum length of the password. The minimum can be from 8 to 40 characters. New passwords must between the minimum length specified and 40 characters. The default value is 8. The total of the other password strength policy parameters (lowercase, uppercase, digits and punctuation) must be less than or equal to the value of the MinLength parameter.

The following example shows a password strength policy that requires passwords to contain at least 3 uppercase characters, 4 lowercase characters and 2 numeric digits; the minimum length of the password is 9 characters.

```
passwdcfg --set -uppercase 3 -lowercase 4 -digits 2 -minlength 9
```

Setting the Password History Policy

The password history policy prevents users from recycling recently used passwords, and is enforced across all user accounts when users are setting their own passwords. The password history policy is enforced only when a new password is defined.

You specify the number of past password values that are disallowed when setting a new password. Allowable password history values range between 1 and 24. The default value is 1, which means the current password can not be reused. The value 2 indicates that both the current and previous passwords cannot be used (and so on, up to 24 passwords).

This policy does not verify that a new password meets a minimal standard of difference from prior passwords, rather, it only determines whether or not a newly-specified password is identical to one of the specified number (1-24) of previously used passwords.

The password history policy is not enforced when an administrator sets a password for another user; instead, the user's password history is preserved and the password set by the administrator is recorded in the user's password history.

Setting the Password Expiration Policy

The password expiration policy forces expiration of a password after a configurable period of time, and is enforced across all user accounts. A warning that password expiration is approaching is displayed when the user logs in. When a user's password expires, he or she must change the password to complete the authentication process and open a user session. You can specify the number of days prior to password expiration during which warnings will commence. Password expiration does not disable or lock out the account.

Use the following attributes to set the password expiration policy:

- MinPasswordAge

  Specifies the minimum number of days that must elapse before user can change a password. MinPasswordAge values range from 0 to 999. The default value is zero. Setting this parameter to a non-zero value discourages users from rapidly changing a password in order to circumvent the password history setting to select a recently-used password. The MinPasswordAge policy is not enforced when an administrator changes the password for another user.

- MaxPasswordAge

  Specifies the maximum number of days that can elapse before a password must be changed, and is also known as the password expiration period. MaxPasswordAge values in range from 0 to 999. The default value is zero. Setting this parameter to zero disables password expiration.

- Warning

  Specifies the number of days prior to password expiration that a warning about password expiration is displayed. Warning values range from 0 to 999. The default value is 0 days.

Note

When MaxPasswordAge is set to a non-zero value, MinPasswordAge and Warning must be set to a value that is less than or equal to MaxPasswordAge.

Upgrade and Downgrade Considerations

If you are upgrading from a v5.0.x environment to v5.2.0, the existing password databases do not contain the state information that implements password expiration. So, when the password expiration policy is first set after an upgrade to v5.2.0, any user who has not changed their password will have their password expiration period set to the maximum password expiration period. You must explicitly define the password expiration for users who have not performed a password change subsequent to the upgrade.

For example:

- March 1st—Using a v5.0.x Fabric OS release. User A changes her password.
- April 1—Upgrade to v5.2.0
- May 1—User B changes his password.
- June 1—The password configuration parameter MaxPasswordAge is set to 90 days.

User A's password will expire on September 1. User B's password will expire on August 1.

Setting the Account Lockout Policy

The account lockout policy disables a user account when that user exceeds a specified number of failed login attempts, and is enforced across all user accounts. You can configure this policy to keep the account locked until explicit administrative action is taken to unlock it, or the locked account can be automatically unlocked after a specified period. Administrators can unlock a locked account at any time.

A failed login attempt counter is maintained for each user on each switch instance. The counters for all user accounts are reset to zero when the account lockout policy is enabled. The counter for an individual account is reset to zero when the account is unlocked after a LockoutDuration period expires.

Note that the account locked state is distinct from the account-disabled state.

Use the following attributes to set the account lockout policy:

- LockoutThreshold

  Specifies the number of times a user can attempt to login using an incorrect password before the account is locked. The number of failed login attempts is counted from the last successful login. LockoutThreshold values range from 0 to 999, and the default value is 0. Setting the value to 0 disables the lockout mechanism.

- LockoutDuration

  Specifies the time, in minutes, after which a previously locked account is automatically unlocked. LockoutDuration values range from 0 to 99999, and the default value is 30. Setting the value to 0 disables lockout duration, and would require a user to seek administrative action to unlock the account. The lockout duration begins with the first login attempt after the LockoutThreshold has been reached. Subsequent failed login attempts do not extend the lockout period.

Managing Fabric OS Users on the RADIUS Server

All existing Fabric OS mechanisms for managing switch-local user accounts and passwords remain functional when the switch is configured to use RADIUS. Changes made to the switch-local database do not propagate to the RADIUS server, nor do the changes affect any account on the RADIUS server.

Switch to RADIUS Server Interaction

When configured to use RADIUS, the switch acts as a Network Access Server (NAS) and RADIUS client. The switch sends all AAA service requests to the RADIUS server, following the RFC 2865 protocol. The RADIUS server receives the request packet, validates the request and sends responses packet back to the switch.

A switch can be configured to try both RADIUS and local switch authentication.

For chassis-based systems such as the SilkWorm 48000 Director, the switch IP addresses are aliases of the physical Ethernet interfaces on the CP blades. When specifying client IP addresses for the logical switches in such systems, make sure the CP IP addresses are used. For accessing both the active and standby CP, and for the purpose of HA failover, both CP IP addresses of a chassis should be included in the RADIUS server configuration.

Creating Fabric OS User Accounts

With RADIUS servers, set up user accounts by their true network wide identity rather than by the account names created on a Fabric OS switch. Along with each account name, assign appropriate switch access roles.

RADIUS supports all the defined RBAC roles described in Table 3-2 on page 3-2.

Users must enter their assigned RADIUS account name and password when logging in to a switch that has been configured with RADIUS. After the RADIUS server authenticates a user, it responds with the assigned switch role in a *Brocade Vendor-Specific Attribute* (VSA). If the response does not have a VSA role assignment, the “user” role is assigned. If no Administrative Domain is assigned then they are assigned to the default Admin Domain AD0.

The syntax used for assigning VSA-based account switch roles on a RADIUS server is described in Table 3-7.

Table 3-7 Syntax for VSA-based Account Roles

Item	Value	Description
Type	26	1 octet
Length	7 or higher	1 octet, calculated by the server
Vendor ID	1588	4 octet, Brocade's SMI Private Enterprise Code
Vendor type	1	1 octet, Brocade-Auth-Role; valid attributes for the Brocade-Auth-Role are: SwitchAdmin ZoneAdmin FabricAdmin BasicSwitchAdmin Operator User Admin
	2	*Optional:* Specifies the Admin Domain member list. See "RADIUS Configuration and Admin Domains" on page 3-20. Brocade-AVPairs1
	3	Brocade-AVPairs2
	4	Brocade-AVPairs3
	5	Brocade-AVPairs4
Vendor length	2 or higher	1 octet, calculated by server, including vendor-type and vendor-length
Attribute-specific data	ASCII string	multiple octet, maximum 253, indicating the name of assigned role and other supported attribute values such as Admin Domain member list.

Windows 2000 IAS

For example, to configure a Windows 2000 IAS server to use VSA to pass the “Admin” role to the switch in the dial-in profile, the configuration specifies the Vendor code (1588), Vendor-assigned attribute number (1), and attribute value (admin), as shown in the following.

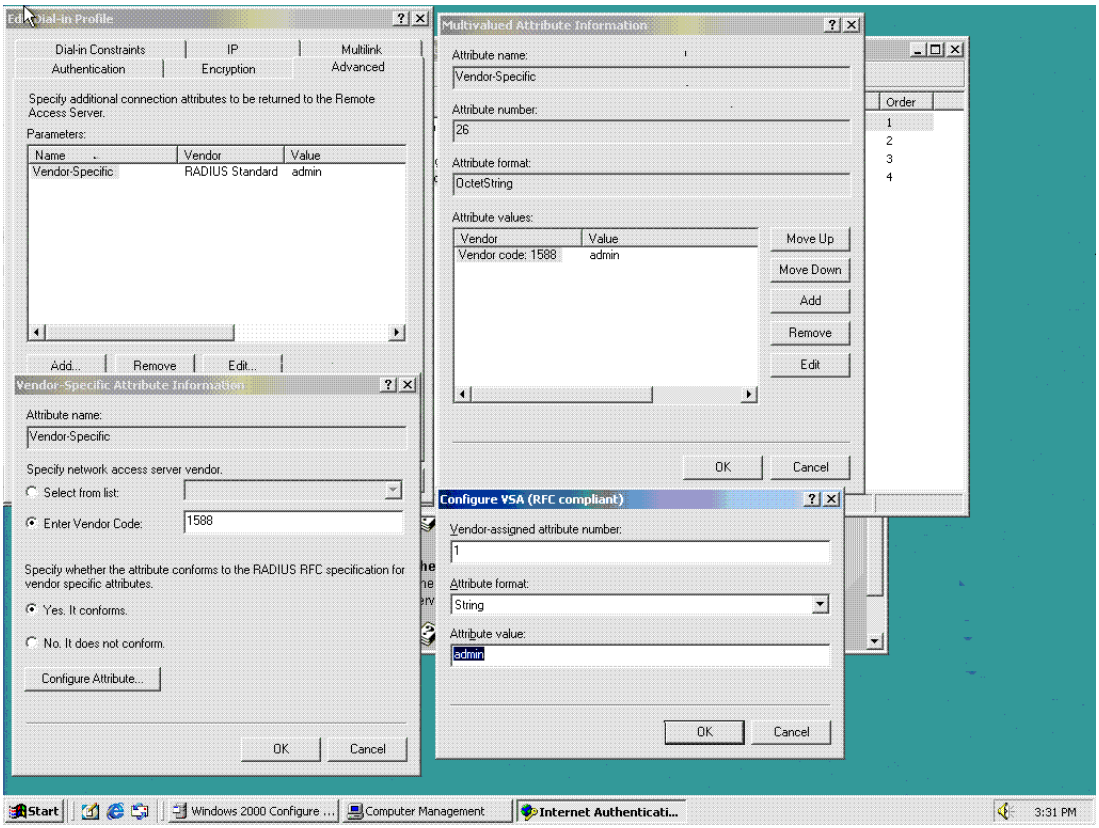

Linux FreeRadius Server

For the configuration on a Linux FreeRadius server, define the following in a vendor dictionary file called dictionary.brocade. Include:

Table 3-8 dictionary.brocade file entries

Include	Key	Value
VENDOR	Brocade	1588
ATTRIBUTE	Brocade-Auth-Role	1 string Brocade
	AdminDomain	

After you have completed the dictionary file, define the role for the user in a configuration file. For example to grant the user jsmith the Admin role, you would add into the configuration file:

```
jsmith  Auth-Type:= Local, User-Password == "jspassword"
        Brocade-Auth-Role = "admin"
```

RADIUS Configuration and Admin Domains

When configuring users with Admin Domains, you must also include the Admin Domain member list. This section describes the way that you configure attribute types for this configuration.

The values for the new attribute types use the syntax *key=val[;key=val]*, where *key* is a text description of attributes, *value* is the attribute value for the given key, = is the separator between key and value, and *;* is an optional separator for multiple key-value pairs.

Multiple key-value pairs can appear for one Vendor-Type code. Key-value pairs with the same key name may get concatenated across multiple Vendor-Type code. You can use any combination of the Vendor-Type codes to specify key-value pairs. Note that a switch always parses these attributes from *Vendor-Type code 2* to V*endor-Type code 4*.

Only two kinds of keys are accepted, *HomeAD* and *ADList*. All other keys are ignored.

- *HomeAD* is the designated home Admin Domain for the account. The valid value is between 0 to 255, inclusive. The first valid HomeAD key-value pair is accepted by the switch, and any additional HomeAD key-value pairs are ignored.
- *ADList* is a comma-separated list of Administrative Domain numbers to which this account is a member. Valid numbers are between 0 to 255, inclusive. A dash between two numbers specifies a range. Multiple ADList key-value pairs within the same or across the different Vendor-Type code are concatenated. Multiple occurrence of the same AD number is ignored.

RADIUS authentication requires that the account have a valid role through the attribute type *Brocade-Auth-Role*. The additional attribute values ADList and HomeAD are optional. If they are unspecified, the account can login with AD0 as its member list and home Admin Domain. If there is an error in ADList or HomeAD specification, the account cannot login until the AD list is corrected; an error message is displayed.

For example, on a Linux FreeRadius Server, the user (user-za) with the following settings takes the “ZoneAdmin” role, with AD member list: 1, 2, 4, 5, 6, 7, 8, 9, 12; the Home Admin Domain will be 1.

```
user-za Auth-Type := Local, User-Password == "password"
Brocade-Auth-Role = "ZoneAdmin",
Brocade-AVPairs1 = "ADList=1,2,6,
Brocade-AVPairs2 = "ADList=4-8;ADList=7,9,12”
```

In the next example, on a Linux FreeRadius Server, the user takes the “Operator” role, with ADList 0 and homeAD 0.

```
user-opr   Auth-Type := Local, User-Password == "password"
Brocade-Auth-Role = "operator",
Brocade-AVPairs1 = "ADList=1,2;HomeAD=2",
Brocade-AVPairs2 = "ADList=-4-8,20;ADList=7,9,12",
```

Setting Up RADIUS AAA Service

Fabric OS v3.2, v4.4.0, and later support RADIUS authentication, authorization, and accounting service (AAA). When configured for RADIUS, a switch becomes a RADIUS client. In this configuration, authentication records are stored in the RADIUS host server database. Login and logout account name, assigned role, and time-accounting records are also stored on the RADIUS server for each user.

By default, RADIUS service is disabled, so AAA services default to the switch local database.

To enable RADIUS service, it is strongly recommended that you access the CLI through an SSH connection so that the shared secret is protected. Multiple login sessions can configure simultaneously, and the last session to apply a change leaves its configuration in effect. After a configuration is applied, it persists after a reboot or an HA failover.

The configuration is chassis-based, so it applies to all logical switches (domains) on the switch and replicates itself on a standby CP blade, if one is present. It is saved in a configuration upload and applied in a configuration download.

You should configure at least two RADIUS servers so that if one fails, the other will assume service. You can set the configuration with both RADIUS service and local authentication enabled so that if all RADIUS servers do not respond (because of power failure or network problems), the switch uses local authentication.

Consider the following effects of the use of RADIUS service on other Fabric OS features:

- When RADIUS service is enabled, all account passwords must be managed on the RADIUS server. The Fabric OS mechanisms for changing switch passwords remain functional; however, such changes affect only the involved switches locally. They do not propagate to the RADIUS server, nor do they affect any account on the RADIUS server.

  When RADIUS is set up for a fabric that contains a mix of switches with and without RADIUS support, the way a switch authenticates users depends on whether or not a RADIUS server is set up for that switch. For a switch with RADIUS support and configuration, authentication bypasses the local password database. For a switch without RADIUS support or configuration, authentication uses the switch's local account names and passwords.

- When Secure Fabric OS secure mode is enabled, the following behaviors apply:
  - Account passwords stored in the switch-local password database are distributed among all switches in the same fabric. RADIUS configuration is not affected.
  - There are separate Admin and Nonfcsadmin roles in secure mode. A nonfcsadmin account on a RADIUS server cannot access FCS switches, even if the account is properly authenticated.
  - If a nonfcsadmin account on a RADIUS server logs in to a switch in nonsecure mode, the switch grants the user Admin role privileges.
- The following behaviors apply to Web Tools:
  - Web Tools client and server keep a session open after a user is authenticated. A password change on a switch invalidates an open session and requires the user to log in again. When integrated with RADIUS, a switch password change on the RADIUS server does not invalidate an existing open session, although a password change on the local switch does.
  - If you cannot log in because of a RADIUS server connection problem, Web Tools displays a message indicating server outage.

Configuring the RADIUS Server

You must know the switch IP address or name to connect to switches. Use the **ipAddrShow** command to display a switch IP address.

For SilkWorm directors (chassis-based systems), the switch IP addresses are aliases of the physical Ethernet interfaces on the CP blades. When specifying client IP addresses for the logical switches in such systems, make sure that the CP blade IP addresses are used. For accessing both the active and standby CP blade, and for the purpose of HA failover, both of the CP blade IP addresses should be included in the RADIUS server configuration.

User accounts should be set up by their true network-wide identity, rather than by the account names created on a Fabric OS switch. Along with each account name, the administrator should assign appropriate switch access roles. To manage a nonsecure fabric, these roles can be User or Admin. To manage a secure fabric, these roles can be User, Admin, or Nonfcsadmin.

When they log in to a switch configured with RADIUS, users enter their assigned RADIUS account names and passwords at the prompt. After RADIUS server authenticates a user, it responds with the assigned switch role in a Brocade Vendor-Specific Attribute (VSA), as defined in the RFC. An Authentication-Accept response without such VSA role assignment automatically assigns the user role.

The following sections describe how to configure a RADIUS server to support Brocade clients under different operating systems.

Linux

The following procedures work for FreeRADIUS on Solaris and Red Hat Linux. FreeRADIUS is a freeware RADIUS server that you can find at:

www.freeradius.org

Follow the installation instructions at the Web site. FreeRADIUS runs on Linux (all versions), FreeBSD, NetBSD, and Solaris. If you make a change to any of the files used in this configuration, you must stop the server and restart it for the changes to take effect.

FreeRADIUS installation places the configuration files in *$PREFIX/etc/raddb*. By default, the PREFIX is */usr/local*.

Configuring RADIUS service on Linux consists of the following tasks:

- Adding the Brocade attribute to the server
- Creating the user
- Enabling clients

To add the Brocade attribute to the server

1. Create and save the file *$PREFIX/etc/raddb/dictionary.brocade* with the following information:

```
#
# Brocade FabricOS v5.0.1 dictionary
#
VENDOR    Brocade            1588
#
# attribute 1 defined to be Brocade-Auth-Role
# string defined in user configuration
#
ATTRIBUTE Brocade-Auth-Role  1      string      Brocade
```

   This defines the Brocade vendor ID as 1588, the Brocade attribute 1 as Brocade-Auth-Role, and it is a string value.

2. Open the file *$PREFIX/etc/raddb/dictionary* in a text editor and add the line:

```
$INCLUDE dictionary.brocade
```

   As a result, the file *dictionary.brocade* is located in the RADIUS configuration directory., and loaded for use by the RADIUS server.

To create the user

1. Open the *$PREFIX/etc/raddb/user* file in a text editor and add user names and roles for users who will be accessing the switch and authenticating RADIUS.

   The user will log in using the role specified with Brocade-Auth-Role. The valid roles include Root, Factory, Admin, SwitchAdmin, and User. The new roles appearing in v5.2.0 are also valid. You must use quotation marks around “password” and “role”.

   For example, to set up an account called JohnDoe with the Admin role:

```
JohnDoe Auth-Type := Local, User-Password == “johnPassword” Brocade-Auth-Role =
“admin”
```

   The next example uses the local system password file to authenticate users.

```
JohnDoe Auth-Type := System, Brocade-Auth-Role = “admin”
```

   When you use NIS for authentication, the only way to enable authentication with the password file is to force the Brocade switch to authenticate using PAP; this requires the *-a pap* option with the **aaaConfig** command.

Clients are the switches that will be using the RADIUS server; each client must be defined. By default, all IP addresses are blocked.

On dual-CP switches (SilkWorm 24000 and 48000), the switch sends its RADIUS request using the IP address of the active CP. When adding clients, add both the active and standby CP IP addresses so that users can still log in the event of a failover.

To enable clients

1. Open the *$PREFIX/etc/raddb/client.config* file in a text editor and add the switches that are to be configured as RADIUS clients.

   For example, to configure the switch at IP address 10.32.170.59 as a client:

```
client 10.32.170.59
   secret      = Secret
   shortname   = Testing Switch
   nastype     = other
```

   In this example, *shortname* is an alias used to easily identify the client. *Secret* is the shared secret between the client and server. Make sure that the shared secret matches that configured on the switch (see "To add a RADIUS server to the switch configuration" on page 3-28).

2. Save the file *$PREFIX/etc/raddb/client.config*; then start the RADIUS server as follows:

```
$PREFIX/sbin/radiusd
```

Windows 2000

Configuring RADIUS service on Windows 2000 consists of the following tasks:

- Installing Internet Authentication Service (IAS)

  For more information and instructions on installing IAS, refer to the Microsoft Web site.

- Enabling the Challenge Handshake Authentication Protocol (CHAP)

  If CHAP authentication is required, then Windows must be configured to store passwords with reversible encryption. Reverse password encryption is not the default behavior; it must be enabled.

Note

If a user is configured prior to enabling reverse password encryption, then the user's password is stored and cannot utilize CHAP. To use CHAP, the password must be reentered after encryption is enabled. If the password is not reentered, then CHAP authentication will not work and the user will be unable to authenticate from the switch.

- Configuring a user

  Internet Authentication Service (IAS) is the Microsoft implementation of a RADIUS server and proxy**.** IAS uses the Windows native user database to verify user login credentials; it does not list specific users, but instead lists *user groups*. Each user group should be associated with a specific switch login role. For example, you should configure a user group for root, admin, factory, switchadmin, and user, and then add any users whose logins you want to associate to the appropriate group.

- Configuring the server

To enable CHAP

1. From the Windows Start menu, select **Programs> Administrative Tools> Local Security Policy** to open the Local Security Settings window.
2. In the Local Security Settings window, expand the **Account Policies** folder and select the **Password Policy** folder.

3. From the list of policies in the Password Policy folder, right-click **Store password using reversible encryption for all users in the domain**, and select **Security** from the pop-up menu.
4. An additional Local Security Settings window appears. Click the **Enabled** radio button and then click **OK**.

To configure users

1. From the Windows Start menu, select **Programs > Administrative Tools > Computer Management** to open the Computer Management window.
2. In the Computer Management window, expand the **Local Users and Groups** folder and select the **Groups** folder.
3. Right-click the **Groups** folder and select **New Group** from the pop-up menu.
4. In the New Group window, provide a Name and Description for the group and click **Add**.
5. In the Select Users or Groups window, select the user (who should already have been configured) you want to add to the group and click **Add**.
6. Repeat this for every user you want to add. When you have completed adding all users, click **OK**.
7. In the New Group window, verify that the users you added in step 4 appear in the Members field; then click **Create** to create this group.

   The new groups are created for each login type (admin, switchAdmin, user).

To configure the RADIUS server

1. From the Windows Start menu, select **Programs > Administrative Tools > Internet Authentication Service** to open the Internet Authentication Service window.
2. In the Internet Authentication Service window, right-click the **Clients** folder and select **New Client** from the pop-up menu.

   Note
   A client is the device that uses the RADIUS server; in this case, it is the switch.

3. In the Add Client window, provide the following:

   Friendly name—The friendly name should be an alias that is easily recognizable as the switch to which you are connecting.

   Protocol—Select **RADIUS** as the protocol.

4. In the Add RADIUS Client window, provide the following:

   Client address (IP or DNS)**—**Enter the IP address of the switch.

   Client-Vendor—Select **RADIUS Standard**.

   Shared secret—Provide a password. Shared secret is a password used between the client device and server to prevent IP address spoofing by unwanted clients. Keep your shared secret password in a safe place. You will need to enter this password in the switch configuration.

   After clicking **Finish**, repeat steps 2 through 4 for all switches on which RADIUS authentication will be used.

5. In the Internet Authentication Service window, right-click the Remote Access Policies folder; then select **New Remote Access Policy** from the pop-up window.

   A remote access policy must be created for each Brocade login role (Root, Admin, Factory, SwitchAdmin, and User) for which you want to use RADIUS. Apply this policy to the user groups that you already created.

6. In the Add Remote Access Policy window, enter an easily identifiable **Policy friendly name** that will enable you to see the switch login for which the policy is being created; then click **Next**.
7. After the Add Remote Access Policy window refreshes, click **Add**.
8. In the Select Attribute window, select **Windows Groups** and click **Add**.
9. In the Groups window, click **Add**.
10. In the Select Groups window, select the user-defined group for which you are creating a policy and click **Add**. After adding all appropriate groups, click **OK**. In the Groups window, click **OK**.
11. In the Add Remote Access Policy window, confirm that the Conditions section displays the group(s) that you selected and click **Next**.

    After the Add Remote Access Policy window refreshes, select the **Grant remote access permission** radio button and click **Next**.

    After the Add Remote Access Policy window refreshes again, click **Edit Profile**.

12. In the Edit Dial-in Profile window, click the **Authentication** tab and check only the **Encrypted Authentication (CHAP)** and **Unencrypted Authentication (PAP, SPAP)** checkboxes; then click the **Advanced** tab and click **Add**.
13. In the Add Attributes window, select **Vendor-Specific** and click **Add**.
14. In the Multivalued Attribute Information window, click **Add**.
15. In the Vendor-Specific Attribute Information window, click the **Enter Vendor Code** radio button and enter the value **1588**. Click the **Yes. It conforms** radio button, and then click **Configure Attribute...**.
16. In the Configure VSA (RFC compliant) window, enter the following values and click **OK**.

    Vendor-assigned attribute number—Enter the value **1**.

    Attribute format—Enter **String**.

    Attribute value—Enter the login role (Root, Admin, Factory, SwitchAdmin, or User) the user group must use to log in to the switch.

17. In the Multivalued Attribute Information window, click **OK**.
18. In the Edit Dial-in Profile window, remove all additional parameters (except the one you just added, “Vendor-Specific”) and click **OK**.
19. In the Add Remote Access Policy window, click **Finish**.
20. After returning to the Internet Authentication Service window, repeat steps 5 through 19 to add additional policies for all Brocade login types you want to use the RADIUS server. After this is done, you can configure the switch.

Configuring RADIUS Servers on the Switch

RADIUS configuration of the switch is controlled by the **aaaConfig** command.

Note

On dual-CP switches (SilkWorm 24000 and 48000), the switch sends its RADIUS request using the IP address of the active CP. When adding clients, add both the active and standby CP IP addresses so that users can still log in the event of a failover.

The following procedures show how to use the **aaaConfig** command to set up a switch for RADIUS service.

RADIUS configuration is chassis-based configuration data. On platforms containing multiple switch instances, the configuration applies to all instances. The configuration is persistent across reboot and firmwareDownload. On a chassis-based system, the command must replicate the configuration to the standby CP.

Multiple login sessions can invoke the command simultaneously. The last session that applies the change be the one whose configuration is in effect. This configuration is persistent after an HA failover.

To display the current RADIUS configuration

1. Connect to the switch and log in as admin.
2. Enter this command:

```
switch:admin> aaaConfig --show
```

If a configuration exists, its parameters are displayed. If RADIUS service is not configured, only the parameter heading line is displayed. Parameters include:

Position	The order in which servers are contacted to provide service
Server	The server names or IP addresses
Port	The server ports
Secret	The shared secrets
Timeouts	The length of time servers have to respond before the next server is contacted
Authentication	The type of authentication being used on servers

To add a RADIUS server to the switch configuration

1. Connect to the switch and log in as admin.
2. Enter this command:

```
switch:admin> aaaConfig --add server [-p port] [-s secret] [-t timeout]
[-a pap | chap]
```

server	Enter either a server name or IP address. Avoid duplicating server listings (that is, listing the same server once by name and again by IP address). Up to five servers can be added to the configuration.
-p *port*	Optionally, enter a server port. The default is port 1812.
-s *secret*	Optionally, enter a shared secret. The default is “sharedsecret”. Secrets can be from 8 to 40 alphanumeric characters long. Make sure that the secret matches that configured on the server.
-t *timeout*	Optionally, enter the length of time (in seconds) that the server has to respond before the next server is contacted. The default is three seconds. Time-out values can range from 1 to 30 seconds.
-a[pap|chap]	Specify PAP or CHAP as authentication protocol.

At least one RADIUS server must be configured before you can enable RADIUS service.

If no RADIUS configuration exists, turning it on triggers an error message. When the command succeeds, the event log indicates that the configuration is enabled or disabled.

Note
When the RADIUS authentication mode is set to **radiuslocal**, you cannot downgrade Fabric OS to any version lower than v5.2.0: previous versions do not support the **radiuslocal** mode.

To enable and disable a RADIUS server

1. Connect to the switch and log in as admin.
2. Enter this command to enable RADIUS + local:

```
switch:admin> aaaconfig --radiuslocal
```

Local is used if the user authentication fails on the RADIUS server.

or to enable RADIUS + localbackup:

```
switch:admin> aaaconfig --radiuslocalbackup
```

.Local is used if the RADIUS servers are not accessible.

To delete a RADIUS server from the configuration

1. Connect to the switch and log in as admin.
2. Enter this command:

```
switch:admin> aaaConfig --remove server | all
```

server	Servers are listed by either name or IP address. Enter either the name or IP address of the server to be removed.

3. At the prompt, enter **y** to complete the command.

When the command succeeds, the event log indicates that the server is removed.

To change a RADIUS server configuration

1. Connect to the switch and log in as admin.
2. Enter this command:

```
switch:admin> aaaConfig --change server [-p port] [-s secret] [-t timeout] [-a
pap | chap]
```

server	Servers are listed by either name or IP address. Enter either the name or IP address of the server to be changed.
-p *port*	Optionally, enter a server port.
-s *secret*	Optionally, enter a shared secret.
-t *timeout*	Optionally, enter the length of time (in seconds) the server has to respond before the next server is contacted.
-a[pap|chap]	Specify PAP or CHAP as authentication protocol.

To change the order in which RADIUS servers are contacted for service

1. Connect to the switch and log in as admin.
2. Enter this command:

```
switch:admin> aaaConfig --move server to_position
```

server	Servers are listed by either name or IP address. Enter either the name or IP address of the server whose position is to be changed.
to_position	Enter the position number to which the server is to be moved.

When the command succeeds, the event log indicates that a server configuration is changed.

Enabling and Disabling Local Authentication as Backup

It is useful to enable local authentication so that the switch can take over authentication locally if the RADIUS servers fail to respond because of power outage or network problems. To enable or disable local authentication, enter the appropriate command:

```
switch:admin> aaaConfig –radiuslocalbackup
```

For details about this command and how it is different from **aaaConfig –radiuslocal**, see Table 3-5 on page 3-6.

When local authentication is enabled and RADIUS servers fail to respond, you can log in to the default switch accounts (admin and user) or any user-defined account. You must know the passwords of these accounts.

When the command succeeds, the event log indicates that local database authentication is disabled or enabled.

Setting the Boot PROM Password

The boot PROM password provides an additional layer of security by protecting the boot PROM from unauthorized use. Setting a recovery string for the boot PROM password enables you to recover a lost boot PROM password by contacting your switch service provider. Without the recovery string, a lost boot PROM password cannot be recovered.

You should set the boot PROM password and the recovery string on all switches, as described in "With a Recovery String" on page 3-30. If your site procedures dictate that you set the boot PROM password without the recovery string, see "Without a Recovery String" on page 3-33.

With a Recovery String

To set the boot PROM password with a recovery string, refer to the section that applies to your switch model.

Note

Setting the boot PROM password requires accessing the boot prompt, which stops traffic flow through the switch until the switch is rebooted. You should perform this procedure during a planned down time.

SilkWorm 200E, 3250, 3850, 3900, 4100, 4900, and 7500

To set the boot PROM password for a switch with a recovery string

1. Connect to the serial port interface as described in "To connect through the serial port" on page 2-3.
2. Reboot the switch.

3. Press **ESC** within four seconds after the message “Press escape within 4 seconds...” displays.

   The following options are available:

Option		Description
1	Start system.	Continues the system boot process.
2	Recovery password.	Lets you set the recovery string and the boot PROM password.
3	Enter command shell.	Provides access to boot parameters.

4. Enter **2.**

   If no password was previously set, the following message displays:

```
Recovery password is NOT set. Please set it now.
```

   If a password was previously set, the following messages display:

```
    Send the following string to Customer Support for password recovery:
    afHTpyLsDo1Pz0Pk5GzhIw==
    Enter the supplied recovery password.
    Recovery Password:
```

5. Enter the recovery password (string).

   The recovery string must be between 8 and 40 alphanumeric characters. A random string that is 15 characters or longer is recommended for higher security. The firmware prompts for this password only once. It is not necessary to remember the recovery string because it is displayed the next time you enter the command shell.

   The following prompt displays:

```
   New password:
```

6. Enter the boot PROM password; then reenter it when prompted. The password must be eight alphanumeric characters (any additional characters are not recorded). Record this password for future use.

   The new password is automatically saved.

7. Reboot the switch.

SilkWorm 24000 and 48000

The boot PROM and recovery passwords must be set for each CP blade on SilkWorm 24000 and 48000 directors:

To set the boot PROM password for a director with a recovery string

1. Connect to the serial port interface on the standby CP blade, as described in “To connect through the serial port” on page 2-3.
2. Connect to the active CP blade by serial or telnet and enter the **haDisable** command to prevent failover during the remaining steps.

3. **SilkWorm 24000 and 48000**: Reboot the standby CP blade by sliding the On/Off switch on the ejector handle of the standby CP blade to Off, and then back to On.
4. Press **ESC** within four seconds after the message “Press escape within 4 seconds...” displays.

   The following options are available:

Option		Description
1	Start system.	Continues the system boot process.
2	Recovery password.	Lets you set the recovery string and the boot PROM password.
3	Enter command shell.	Provides access to boot parameters.

5. Enter **2**.

   If no password was previously set, the following message displays:

```
Recovery password is NOT set. Please set it now.
```

   If a password was previously set, the following messages display:

```
    Send the following string to Customer Support for password recovery:
    afHTpyLsDo1Pz0Pk5GzhIw==
    Enter the supplied recovery password.
    Recovery Password:
```

6. Enter the recovery password (string).

   The recovery string must be between 8 and 40 alphanumeric characters. A random string that is 15 characters or longer is recommended for higher security. The firmware only prompts for this password once. It is not necessary to remember the recovery string because it is displayed the next time you enter the command shell.

   The following prompt displays:

```
New password:
```

7. Enter the boot PROM password; then reenter it when prompted. The password must be eight alphanumeric characters (any additional characters are not recorded). Record this password for future use.

   The new password is automatically saved (the **saveEnv** command is not required).
8. Connect to the active CP blade by serial or telnet and enter the **haEnable** command to restore high availability; then fail over the active CP blade by entering the **haFailover** command.

   Traffic flow through the active CP blade resumes when the failover is complete.
9. Connect the serial cable to the serial port on the new standby CP blade (previously the active CP blade).
10. Repeat step 2 through step 7 for the new standby CP blade (each CP blade has a separate boot PROM password).
11. Connect to the active CP blade by serial or telnet and enter the **haEnable** command to restore high availability.

Without a Recovery String

Although you can set the boot PROM password without also setting the recovery string, it is strongly recommended that you set both the password and the string as described in "With a Recovery String" on page 3-30. If your site procedures dictate that you must set the boot PROM password without the string, follow the procedure that applies to your switch model.

Note

Setting the boot PROM password requires accessing the boot prompt, which stops traffic flow through the switch until the switch is rebooted. You should perform this procedure during a planned down time.

SilkWorm 200E, 3250, 3850, 3900, 4100, 4900, and 7500

To set the boot PROM password for a switch without a recovery string

1. Create a serial connection to the switch as described in "To connect through the serial port" on page 2-3.
2. Reboot the switch by entering the **reboot** command.
3. Press **ESC** within four seconds after the message "Press escape within 4 seconds..." displays.

   The following options are available:

Option		Description
1	Start system.	Continues the system boot process.
2	Recovery password.	Lets you set the recovery string and the boot PROM password.
3	Enter command shell.	Provides access to boot parameters.

4. Enter **3**.
5. At the shell prompt, enter the **passwd** command.

   Note

   The **passwd** command only applies to the boot PROM password when it is entered from the boot interface.

6. Enter the boot PROM password at the prompt, then reenter it when prompted. The password must be eight alphanumeric characters (any additional characters are not recorded). Record this password for future use.
7. Enter the **saveEnv** command to save the new password.
8. Reboot the switch by entering the **reset** command.

SilkWorm 24000 and 48000

On the SilkWorm 24000 and 48000, set the password on the standby CP blade, fail over, and then set the password on the previously active (now standby) CP blade to minimize disruption to the fabric:

To set the boot PROM password for a director without a recovery string

1. Determine the active CP blade by opening a telnet session to either CP blade, connecting as admin, and entering the **haShow** command.
2. Connect to the active CP blade by serial or telnet and enter the **haDisable** command to prevent failover during the remaining steps.
3. Create a serial connection to the standby CP blade as described in "To connect through the serial port" on page 2-3.
4. **SilkWorm 24000 and 48000**: Reboot the standby CP blade by sliding the On/Off switch on the ejector handle of the standby CP blade to Off, and then back to On.

   This causes the blade to reset.
5. Press **ESC** within four seconds after the message "Press escape within 4 seconds..." displays.

   The following options are available:

Option		Description
1	Start system.	Continues the system boot process.
2	Recovery password.	Lets you set the recovery string and the boot PROM password.
3	Enter command shell.	Provides access to boot parameters.

6. Enter **3**.
7. Enter the **passwd** command at the shell prompt.

Note

The **passwd** command applies only to the boot PROM password when it is entered from the boot interface.

8. Enter the boot PROM password at the prompt; then reenter it when prompted. The password must be eight alphanumeric characters (any additional characters are not recorded). Record this password for future use.
9. Enter the **saveEnv** command to save the new password.
10. Reboot the standby CP blade by entering the **reset** command.
11. Connect to the active CP blade by serial or telnet and enter the **haEnable** command to restore high availability; then fail over the active CP blade by entering the **haFailover** command.

    Traffic resumes flowing through the newly active CP blade after it has completed rebooting.
12. Connect the serial cable to the serial port on the new standby CP blade (previously the active CP blade).
13. Repeat step 3 through step 10 for the new standby CP blade.
14. Connect to the active CP blade by serial or telnet and enter the **haEnable** command to restore high availability.

Recovering Forgotten Passwords

If you know the root password, you can use this procedure to recover the user, admin, and factory passwords.

To recover passwords

1. Open a CLI connection (serial or telnet) to the switch. If secure mode is enabled, connect to the primary FCS switch.
2. Log in as root.
3. Enter the command for the type of password that was lost:

```
passwd user
passwd admin
passwd factory
```

4. Enter the requested information at the prompts.

To recover a lost root password, contact your switch service provider.

To recover a lost boot PROM password, contact your switch service provider. You must have previously set a recovery string to recover the boot PROM password.

Configuring Standard Security Features

Chapter 4

This chapter provides information and procedures for configuring standard Fabric OS security features such as account and password management. Additional security features are available when secure mode is enabled. For information about licensed security features available in Secure Fabric OS, see the *Secure Fabric OS Administrator's Guide*.

- "Secure Protocols" on page 4-1
- "Ensuring Network Security" on page 4-3
- "Configuring the Telnet Interface" on page 4-4
- "Blocking Listeners" on page 4-5
- "Accessing Switches and Fabrics" on page 4-5
- "Configuring for the SSL Protocol" on page 4-7
- "Configuring for SNMP" on page 4-14
- "Configuring Secure File Copy" on page 4-23

Secure Protocols

Fabric OS supports the secure protocols shown in Table 4-1.,

Table 4-1 Secure Protocol Support

Protocol	Description
SSL	Supports SSLv3, 128-bit encryption by default. Fabric OS uses SSL to support HTTPS. A certificate must be generated and installed on each switch to enable SSL.
HTTPS	Web Tools supports the use of HTTPS.
Secure File Copy (scp)	Configuration upload and download support the use of scp.
SNMPv3	SNMPv1 is also supported.

Simple Network Management Protocol (SNMP) is a standard method for monitoring and managing network devices. Using SNMP components, you can program tools to view, browse, and manipulate Brocade switch variables and set up enterprise-level management processes.

Every Brocade switch carries an SNMP agent and management information base (MIB). The agent accesses MIB information about a device and makes it available to a network manager station. You can manipulate information of your choice by *trapping* MIB elements using the Fabric OS CLI, Web Tools, or Fabric Manager.

The SNMP Access Control List (ACL) provides a way for the administrator to restrict SNMP get/set operations to certain hosts/IP addresses. This is used for enhanced management security in the storage area network.

For details on Brocade MIB files, naming conventions, loading instructions, and information about using Brocade's SNMP agent, see the *Fabric OS MIB Reference*.

Table 4-2 describes additional software or certificates that you must obtain to deploy secure protocols.

Table 4-2 Items Needed to Deploy Secure Protocols

Protocol	Host Side	Switch Side
Secure telnet (sectelnet)	Sectelnet client	License not required, but a switch certificate issued by Brocade is required
SSH	SSH client	None
HTTPS	No requirement on host side except a browser that supports HTTPS	Switch IP certificate for SSL
Secure File Copy (scp)	SSH daemon, scp server	None
SNMPv3, SNMPv1	None	None

The security protocols are designed with the four main usage cases described in Table 4-3.

Table 4-3 Main Security Scenarios

Fabric	Management Interfaces	Comments
Nonsecure	Nonsecure	No special setup is need to use telnet or HTTP. A Brocade switch certificate must be installed if sectelnet is used.
Nonsecure	Secure	Secure protocols might be used. An SSL switch certificate must be installed if SSH/HTTPS is used.
Secure	Secure	Secure protocols are supported on Fabric OS v4.4.0 (and later) switches. Switches running earlier Fabric OS versions can be part of the secure fabric, but they do not support secure management. Secure management protocols must be configured for each participating switch. Nonsecure protocols might be disabled on nonparticipating switches. If SSL is used, then certificates must be installed.

Table 4-3 Main Security Scenarios (Continued)

Fabric	Management Interfaces	Comments
Secure	Nonsecure	You must use sectelnet because telnet is not allowed in secure mode. Nonsecure management protocols are necessary under these circumstances: • The fabric contains switches running Fabric OS v3.2.0. • The presence of software tools that do not support Secure protocols: for example, Fabric Manager v4.0.0. • The fabric contains switches running Fabric OS versions earlier than v4.4.0. Nonsecure management is enabled by default.

Ensuring Network Security

To ensure security, Fabric OS supports secure shell (SSH) encrypted sessions in v4.1.x and later. SSH encrypts all messages, including the client's transmission of password during login. The SSH package contains a daemon (sshd), which runs on the switch. The daemon supports a wide variety of encryption algorithms, such as Blowfish-CBC and AES.

Note

To maintain a secure network, you should avoid using telnet (you can use secTelnet if you are using Fabric OS v2.6 or later, or v3.1 or later) or any other unprotected application when you are working on the switch. For example, if you use telnet to connect to a machine, and then start an SSH or secure telnet session from that machine to the switch, the communication to the switch is in clear text and therefore is not secure.

The FTP protocol is also not secure. When you use FTP to copy files to or from the switch, the contents are in clear text. This includes the remote FTP server's login and password. This limitation affects the following commands: **saveCore**, **configUpload**, **configDownload**, and **firmwareDownload**.

Commands that require a secure login channel must be issued from an original SSH session. If you start an SSH session, and then use the login command to start a nested SSH session, commands that require a secure channel will be rejected.

Fabric OS v4.1.0 and later supports SSH protocol v2.0 (ssh2). For more information on SSH, refer to the SSH IETF Web site:

http://www.ietf.org/ids.by.wg/secsh.html

Refer to *SSH, The Secure Shell: The Definitive Guide* by Daniel J. Barrett, Richard Silverman.

Fabric OS v4.4.0 and later comes with the SSH server preinstalled; however, you must select and install the SSH client. For information on installing and configuring the F-Secure SSH client, refer to the Web site:

http://www.f-secure.com

Configuring the Telnet Interface

Telnet is enabled by default. To prevent users from passing clear text passwords over the network when they connect to the switch, you can disable the telnet interface.

Note

Before disabling the telnet interface, make sure you have an alternate method of establishing a connection with the switch.

To disable telnet

1. Connect to the switch and log in as admin.

   Connect through some other means than telnet: for example, through SSH.

2. Enter the following command:

```
configure telnetd
```

3. In response to the System Services prompt, enter **y**.
4. In response to the telnetd prompt, enter **off**.

   The telnet interface is disabled. If you entered the command during a standard telnet session, all telnet sessions terminate.

```
switch:admin> configure telnetd
Not all options will be available on an enabled switch.
To disable the switch, use the "switchDisable" command.
Configure...
 ssl attributes (yes, y, no, n): [no]
  http attributes (yes, y, no, n): [no]
  snmp attributes (yes, y, no, n): [no]
  rpcd attributes (yes, y, no, n): [no]
  cfgload attributes (yes, y, no, n): [no]

[31454]: Read 1 license entries for generation 1.
[31454]: Read 1 license records.
  System services (yes, y, no, n): [no] y

    rstatd (on, off): [off]
    rusersd (on, off): [off]
    telnetd (on, off): [on] off
```

To enable telnet

1. Connect to the switch through a means other than telnet (for example, SSH) and log in as admin.
2. Enter the following command:

```
switch:admin> configure telnetd
```

3. In response to the System Services prompt, enter **y**.
4. In response to the telnetd prompt, enter **on**.

   The telnet interface is enabled.

Blocking Listeners

SilkWorm switches block Linux subsystem listener applications that are not used to implement supported features and capabilities. Table 4-4 lists the listener applications that SilkWorm switches either block or do not start.

Table 4-4 Blocked Listener Applications

Listener Application	SilkWorm 24000 and 48000 Directors	SilkWorm 200E, 3250, 3850, 3900, 4100, 4900 and 7500 Switches, FC4-48 and FC4-16IP blades
chargen	Do not start	Do not start
echo	Do not start	Do not start
daytime	Do not start	Do not start
discard	Do not start	Do not start
ftp	Do not start	Do not start
rexec	Block with packet filter	Do not start
rsh	Block with packet filter	Do not start
rlogin	Block with packet filter	Do not start
time	Block with packet filter	Do not start
rstats	Do not start	Do not start
rusers	Do not start	Do not start

Accessing Switches and Fabrics

If you are using the FC-FC Routing Service, be aware that you cannot execute the **secModeEnable** command on backbone fabrics and you cannot run this command in secure mode when a backbone fabric is connected to edge fabrics. See Chapter 11, “Using the FC-FC Routing Service” for details about the FC-FC Routing Service and it relationship with Secure Fabric OS.

Table 4-5 lists the defaults for accessing hosts, devices, switches, and zones.

Table 4-5 Access Defaults

	Access Default
Hosts	Any host can access the fabric by SNMP
	Any host can telnet to any switch in the fabric
	Any host can establish an HTTP connection to any switch in the fabric
	Any host can establish an API connection to any switch in the fabric

Table 4-5 Access Defaults (Continued)

	Access Default
Devices	All device ports can access SES
	All devices can access the management server
	Any device can connect to any FC port in the fabric
Switch access	Any switch can join the fabric
	All switches in the fabric can be accessed through serial port
Zoning	Node WWNs can be used for WWN-based zoning

Port Configuration

Table 4-6 lists the ports used. This table provides the information to make it clearer when configuring the switch, taking into consideration firewalls and other devices that may sit between switches in the fabric or between the managers and the switch.

Table 4-6 Port Information

Port	Type	Common Use	Comment
22	TCP	SSH	
23	TCP	Telnet	Use the **configure** command to disable the telnet service.
37	TCP	NTP	
80	TCP	HTTP	Use the **configure** command to disable the port.
111	TCP	sunrpc	This port is used by Platform API. Use the **configure** command to enable or disable daemons.
161	UDP	SNMP	Disable the SNMP service on the remote host if you do not use it, or filter incoming UDP packets going to this port.
443	TCP	HTTPS	Supported service (Web Tools). Disable this port using the **configure** command.
512	TCP	exec	
513	TCP	login	
514	TCP	shell	
897	TCP		This port is used by the Platform API. Disable this port using the **configure** command.

Configuring for the SSL Protocol

Fabric OS v4.4.0 and later supports secure sockets layer (SSL) protocol, which provides secure access to a fabric through Web-based management tools like Web Tools. SSL support is a standard Fabric OS feature; it is independent of Secure Fabric OS, which requires a license and separate certification.

Switches configured for SSL grant access to management tools through hypertext transfer protocol-secure links (which begin with *https://*) instead of standard links (which begin with *http://*).

SSL uses public key infrastructure (PKI) encryption to protect data transferred over SSL connections. PKI is based on digital certificates obtained from an Internet Certificate Authority (CA), which acts as the trusted key agent.

Certificates are based on the switch IP address or fully qualified domain name (FQDN), depending on the issuing CA. If you change a switch IP address or FQDN after activating an associated certificate, you might have to obtain and install a new certificate. Check with the CA to verify this possibility, and plan these types of changes accordingly.

Browser and Java Support

Fabric OS supports the following Web browsers for SSL connections:

- Internet Explorer (Microsoft Windows)
- Mozilla (Solaris and Red Hat Linux)

In countries that allow the use of 128-bit encryption, you should use the latest version of your browser. For example, Internet Explorer 6.0 and later supports 128-bit encryption by default. You can display the encryption support (called “cipher strength”) using the Internet Explorer **Help:About** menu option. If you are running an earlier version of Internet Explorer, you might be able to download an encryption patch from the Microsoft Web site at *http://www.microsoft.com*.

You should upgrade to the Java 1.4.2_03 Plug-in on your management workstation. To find the Java version that is currently running, open the Java console and look at the first line of the window.

For more details on levels of browser and Java support, see the *Web Tools Administrator’s Guide*.

Summary of SSL Procedures

You configure for SSL by obtaining, installing, and activating digital certificates for SSL support. Certificates are required on all switches that are to be accessed through SSL.

You also need to install a certificate to the Java Plug-in on the management workstation, and you might need to add a certificate to your Web browser.

Configuring for SSL involves these major steps, which are shown in detail in the next sections:

1. Choose a CA.
2. On each switch:
    a. Generate a public/private key (**secCertUtil genkey** command).
    b. Generate a certificate signing request (CSR) (**secCertUtil gencsr** command) and store the CSR on an FTP server (**secCertUtil export** command).

3. Obtain the certificates from the CA.

   You can request a certificate from a CA through a Web browser. After you request a certificate, the CA either sends certificate files by email (public) or gives access to them on a remote host (private). Typically, the CA provides the certificate files listed in Table 4-7.

   Table 4-7 SSL Certificate Files

Certificate File	Description
name.crt	The switch certificate.
*name*Root.crt	The root certificate. Typically, this certificate is already installed in the browser, but if not, you must install it.
*name*CA.crt	The CA certificate. It is not necessary to install this, but you can if you want the CA name to be displayed in the browser window.

4. On each switch, install and then activate the certificate.
5. If necessary, install the root certificate to the browser on the management workstation.
6. Add the root certificate to the Java Plug-in keystore on the management workstation.

Choosing a Certificate Authority

To ease maintenance and allow secure out-of-band communication between switches, consider using one certificate authority (CA) to sign all management certificates for a fabric. If you use different CAs, management services operate correctly, but the Web Tools Fabric Events button is unable to retrieve events for the entire fabric.

Each CA (for example, Verisign or GeoTrust) has slightly different requirements; for example, some generate certificates based on IP address, while others require an FQDN, and most require a 1024-bit public/private key while some might accept a 2048-bit key. Consider your fabric configuration, check CA Web sites for requirements, and gather all the information that the CA requires.

Generating a Public/Private Key

Perform this procedure on each switch:

1. Connect to the switch and log in as admin.
2. Enter this command to generate a public/private key pair:

```
switch:admin> seccertutil genkey
```

   The system reports that this process will disable secure protocols, delete any existing CSR, and delete any existing certificates.

3. Respond to the prompts to continue and select the key size:

```
Continue (yes, y, no, n): [no] y
Select key size [1024 or 2048]: 1024
Generating new rsa public/private key pair
Done.
```

Because CA support for the 2048-bit key size is limited, you should select 1024 in most cases.

Generating and Storing a CSR

After generating a public/private key (see "Generating a Public/Private Key," earlier), perform this procedure on each switch:

1. Connect to the switch and log in as admin.
2. Enter this command:

```
switch:admin> seccertutil gencsr
```

3. Enter the requested information:

```
Country Name (2 letter code, eg, US):US
State or Province Name (full name, eg, California):California
Locality Name (eg, city name):San Jose
Organization Name (eg, company name):Brocade
Organizational Unit Name (eg, department name):Eng
Common Name (Fully qualified Domain Name, or IP address): 192.1.2.3
Generating CSR, file name is: 192.1.2.3.csr
Done.
```

Your CA might require specific codes for Country, State or Province, Locality, Organization, and Organizational Unit names. Make sure that your spelling is correct and matches the CA requirements. If the CA requires that the Common Name be specified as an FQDN, make sure that the fully qualified domain name is set on the domain name server.

4. Enter this command to store the CSR:

```
switch:admin> seccertutil export
```

5. Enter the requested information:

```
Select protocol [ftp or scp]: ftp
Enter IP address: 192.1.2.3
Enter remote directory: path_to_remote_directory
Enter Login Name: your account
Enter Password: your password
Success: exported CSR.
```

If you are set up for secure file copy protocol, you can select it; otherwise, select ftp. Enter the IP address of the switch on which you generated the CSR. Enter the remote directory name of the FTP server to which the CSR is to be sent. Enter your account name and password on the server.

Obtaining Certificates

Check the instructions on the CA Web site; then, perform this procedure for each switch:

1. Generate and store the CSR as described in "Generating and Storing a CSR" on page 4-9.
2. Open a Web browser window on the management workstation and go to the CA Web site. Follow the instructions to request a certificate. Locate the area in the request form into which you are to paste the CSR.
3. Through a telnet window, connect to the switch and log in as admin.
4. Enter this command:

```
switch:admin> seccertutil showcsr
```

   The contents of the CSR is displayed.

5. Locate the section that begins with "BEGIN CERTIFICATE REQUEST" and ends with "END CERTIFICATE REQUEST".
6. Copy and paste this section (including the BEGIN and END lines) into the area provided in the request form; then, follow the instructions to complete and send the request.

It might take several days to receive the certificates. If the certificates arrive by email, save them to an FTP server. If the CA provides access to the certificates on an FTP server, make note of the path name and make sure you have a login name and password on the server.

Installing a Switch Certificate

Perform this procedure on each switch:

1. Connect to the switch and log in as admin.
2. Enter this command:

```
switch:admin> seccertutil import
```

3. Select a protocol, enter the IP address of the host on which the switch certificate is saved, and enter your login name and password:

```
Select protocol [ftp or scp]: ftp
Enter IP address:  192.10.11.12
Enter remote directory: path_to_remote_directory
Enter certificate name (must have ".crt" suffix):192.1.2.3.crt
Enter Login Name: your_account
Enter Password: *****
Success: imported certificate [192.1.2.3.crt].
To use this certificate, run the configure command to activate it
```

   The certificate is downloaded to the switch.

Activating a Switch Certificate

Enter the **configure** command and respond to the prompts that apply to SSL certificates:

SSL attributes	Enter **yes**.
Certificate File	Enter the name of the switch certificate file: for example, **192.1.2.3.crt**.
CA Certificate File	If you want the CA name to be displayed in the browser window, enter the name of the CA certificate file; otherwise, skip this prompt.
Select length of crypto key	Enter the encryption key length (**40**, **56**, or **128**).
HTTP attributes	Enter **yes**.
Secure HTTP enabled	Enter **yes**.

For example:

```
Configure...
System services (yes, y, no, n): [no]
  ssl attributes (yes, y, no, n): [no] yes
Certificate File. (filename or none): [10.33.13.182.crt] 192.1.2.3.crt
        CA Certificate File. (filename or none): [none]
        Select length of crypto key.
          (Valid values are 40, 56, and 128.): (40..128) [128]
http attributes (yes, y, no, n): [no] yes
HTTP Enabled (yes, y, no, n): [yes] no
        Secure HTTP Enabled (yes, y, no, n): [no] yes
```

After you exit the **configure** command, the HTTP daemon restarts automatically to handle HTTPS requests.

Configuring the Browser

The root certificate might already be installed on your browser, but if not, you must install it. To see whether it is already installed, check the certificate store on your browser.

The next procedures are guides for installing root certificates to Internet Explorer and Mozilla browsers. For more detailed instructions, refer to the documentation that came with the certificate.

To check and install root certificates on Internet Explorer

1. From the browser Tools menu, select **Internet Options**.
2. Click the **Content** tab.
3. Click **Certificates**.
4. Click the various tabs and scroll the lists to see if the root certificate is listed. If it is listed, you do not need to install it, forgo the remainder of this procedure.
5. If the certificate is not listed, click **Import**.
6. Follow the instructions in the Certificate Import wizard to import the certificate.

To check and install root certificates on Mozilla

1. From the browser Edit menu, select **Preferences**.
2. In the left pane of the Preferences window, expand the **Privacy & Security** list and select **Certificates**.
3. In the right pane, click **Manage Certificates**.
4. In the next window, click the **Authorities** tab.
5. Scroll the authorities list to see if the root certificate is listed. (For example, its name might have the form *name*Root.crt.) If it is listed, you do not need to install it; forgo the remainder of this procedure.
6. If the certificate is not listed, click **Import**.
7. Browse to the certificate location and select the certificate. (For example, select *name*Root.crt.)
8. Click **Open** and follow the instructions to import the certificate.

Installing a Root Certificate to the Java Plug-in

For information on Java requirements, see "Browser and Java Support" on page 4-7.

This procedure is a guide for installing a root certificate to the Java Plug-in on the management workstation. If the root certificate is not already installed to the plug-in, you should install it. For more detailed instructions, refer to the documentation that came with the certificate and to the Sun Microsystems Web site (*www.sun.com*).

1. Copy the root certificate file from its location on the FTP server to the Java Plug-in bin. For example, the bin location might be:

```
C: \program files\java\j2re1.4.2_03\bin
```

2. Open a Command Prompt window and change directory to the Java Plug-in bin.
3. Enter the **keytool** command and respond to the prompts:

```
C:\Program Files\Java\j2re1.4.2_03\bin> keytool -import -alias  RootCert -file
RootCert.crt -keystore ..\lib\security\RootCerts
Enter keystore password:  changeit
Owner: CN=Brocade, OU=Software, O=Brocade Communications, L=San Jose,
ST=California, C=US
Issuer: CN=Brocade, OU=Software, O=Brocade Communications, L=San Jose,
ST=California, C=US
Serial number: 0
Valid from: Thu Jan 15 16:27:03 PST 2004 until: Sat Feb 14 16:27:03 PST 2004
Certificate fingerprints:
         MD5:  71:E9:27:44:01:30:48:CC:09:4D:11:80:9D:DE:A5:E3
         SHA1: 06:46:C5:A5:C8:6C:93:9C:FE:6A:C0:EC:66:E9:51:C2:DB:E6:4F:A1
Trust this certificate? [no]:  yes
Certificate was added to keystore
```

In the example, **changeit** is the default password and RootCert is an example root certificate name.

Displaying and Deleting Certificates

Table 4-8 summarizes the commands for displaying and deleting certificates. For details on the commands, see the *Fabric OS Command Reference*.

Table 4-8 Commands for Displaying and Deleting SSL Certificates

Command	Description
secCertUtil show	Displays the state of the SSL key and a list of installed certificates
secCertUtil show *filename*	Displays the contents of a specific certificate
secCertUtil showcsr	Displays the contents of a CSR
secCertUtil delete *filename*	Deletes a specified certificate
secCertUtil delcsr	Deletes a CSR

Troubleshooting Certificates

If you receive messages in the browser or in a pop-up window when logging in to the target switch using HTTPS, see Table 4-9.

Table 4-9 SSL Messages and Actions

Message	Action
The page cannot be displayed	The SSL certificate is not installed correctly or HTTPS is not enabled correctly. Make sure that the certificate has not expired, that HTTPS is enabled, and that certificate file names are configured correctly.
The security certificate was issued by a company you have not chosen to trust….	The certificate is not installed in the browser. Install it as described in "Configuring the Browser" on page 4-11.
The security certificate has expired or is not yet valid	Either the certificate file is corrupted or it needs to be updated. Click **View Certificate** to verify the certificate content. If it is corrupted or out of date, obtain and install a new certificate.
The name on the security certificate is invalid or does not match the name of the site file	The certificate is not installed correctly in the Java Plug-in. Install it as described in "Installing a Root Certificate to the Java Plug-in" on page 4-12.
This page contains both secure and nonsecure items. Do you want to display the nonsecure items?	Click **No** in this pop-up window. The session opens with a closed lock icon on the lower-right corner of the browser, indicating an encrypted connection.

Configuring for SNMP

You can configure for the automatic transmission of Simple Network Management Protocol (SNMP) information to management stations. SNMPv3 and SNMPv1 are supported.

The configuration process involves configuring the SNMP agent and configuring SNMP traps. The following commands are used in the process:

- Use the **configure** command to set the security level. You can specify no security, authentication only, or authentication and privacy.
- Use the **snmpConfig** command to configure the SNMP agent and traps for SNMPv3 or SNMPv1 configurations.
- If necessary for backward compatibility, you can use these legacy commands to configure for SNMP v1:
  - Use the **agtCfgShow**, **agtCfgset**, and **agtCfgDefault** commands to configure the SNMPv1 agent.
  - Use the **snmpMibCapSet** command to filter at the trap level and the **snmpMibCapShow** command to display the trap filter values.

The SNMP trap configuration specifies the MIB trap elements to be used to send information to the SNMP management station. There are two main MIB trap choices:

- Brocade-specific MIB trap

  Associated with the Brocade-specific SilkWorm MIB (SW-MIB), this MIB monitors SilkWorm switches specifically.

- FibreAlliance MIB trap

  Associated with the FibreAlliance MIB (FA-MIB), this MIB manages SAN switches and devices from any company that complies with FibreAlliance specifications.

If you use both SW-MIB and FA-MIB, you might receive duplicate information. You can disable the FA-MIB, but not the SW-MIB.

You can also use these additional MIBs and their associated traps:

- FICON-MIB (for FICON environments)
- HA-MIB (for SilkWorm 24000 models)
- SW-EXTTRAP
  Includes the swSsn (Software Serial Number) as a part of Brocade SW traps. It is also used in conjunction with the legacy SilkWorm 6400 integrated fabrics product to provide detailed group information for a particular trap.

For information on Brocade MIBs, see the *Fabric OS MIB Reference*.

For information on the specific commands used in these procedures, see online help or to the *Fabric OS Command Reference*.

Setting the Security Level

Use the **configure** command to set the security level (called “SNMP attributes”). You can specify no security, authentication only, or authentication and privacy. For example, to configure for authentication and privacy:

```
switch:admin> configure

Not all options will be available on an enabled switch.
To disable the switch, use the "switchDisable" command.

Configure...

  System services (yes, y, no, n): [no]
  ssl attributes (yes, y, no, n): [no]
  http attributes (yes, y, no, n): [no]
  snmp attributes (yes, y, no, n): [no] y

         Select SNMP Security Level:
          (0 = No security, 1 = Authentication only, 2 = Authentication and Privacy):
(0..2) [0] 2
```

Using the snmpConfig Command

Use the **snmpConfig --set** command to change either the SNMPv3 or SNMPv1 configuration. You can also change access control, MIB capability, and system group.

Sample SNMPv3 configuration

```
switch:admin> snmpconfig --set snmpv3

     SNMPv3 user configuration:
     User (rw): [snmpadmin1] adminuser
     Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3] 1
     New Auth Passwd:
     Verify Auth Passwd:
     Priv Protocol [DES(1)/noPriv[2]): (1..2) [2] 1
     New Priv Passwd:
     Verify Priv Passwd:
     User (rw): [snmpadmin2] shauser
     Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3] 2
     New Auth Passwd:
     Verify Auth Passwd:
     Priv Protocol [DES(1)/noPriv[2]): (1..2) [2] 1
     New Priv Passwd:
     Verify Priv Passwd:
     User (rw): [snmpadmin3] nosec
     Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
     Priv Protocol [DES(1)/noPriv[2]): (2..2) [2]
     User (ro): [snmpuser1]
     Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (3..3) [3]
     Priv Protocol [DES(1)/noPriv[2]): (2..2) [2]
     User (ro): [snmpuser2]
     Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (3..3) [3]
     Priv Protocol [DES(1)/noPriv[2]): (2..2) [2]
     User (ro): [snmpuser3]
     Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (3..3) [3]
     Priv Protocol [DES(1)/noPriv[2]): (2..2) [2]

     SNMPv3 trap recipient configuration:
     Trap Recipient's IP address in dot notation: [0.0.0.0] 192.168.45.90
     UserIndex: (1..6) [1]
     Trap recipient Severity level : (0..5) [0] 4
     Trap Recipient's IP address in dot notation: [0.0.0.0] 192.168.45.92
     UserIndex: (1..6) [2]
     Trap recipient Severity level : (0..5) [0] 2
     Trap Recipient's IP address in dot notation: [0.0.0.0]
     Trap Recipient's IP address in dot notation: [0.0.0.0]
     Trap Recipient's IP address in dot notation: [0.0.0.0]
     Trap Recipient's IP address in dot notation: [0.0.0.0]
     Committing configuration...done.
```

Sample SNMPv1 configuration

```
switch:admin> snmpconfig --set snmpv1

SNMP community and trap recipient configuration:
Community (rw): [Secret C0de] admin
Trap Recipient's IP address in dot notation: [0.0.0.0] 10.32.225.1
Trap recipient Severity level : (0..5) [0] 1
Community (rw): [OrigEquipMfr]
Trap Recipient's IP address in dot notation: [10.32.225.2]
Trap recipient Severity level : (0..5) [1]
Community (rw): [private]
Trap Recipient's IP address in dot notation: [10.32.225.3]
Trap recipient Severity level : (0..5) [2]
Community (ro): [public]
Trap Recipient's IP address in dot notation: [10.32.225.4]
Trap recipient Severity level : (0..5) [3]
Community (ro): [common]
Trap Recipient's IP address in dot notation: [10.32.225.5]
Trap recipient Severity level : (0..5) [4]
Community (ro): [FibreChannel]
Trap Recipient's IP address in dot notation: [10.32.225.6]
Trap recipient Severity level : (0..5) [5]
Committing configuration...done.
```

Sample accessControl configuration

```
switch:admin> snmpconfig --set accessControl

     SNMP access list configuration:
     Access host subnet area in dot notation: [0.0.0.0] 192.168.0.0
     Read/Write? (true, t, false, f): [true]
     Access host subnet area in dot notation: [0.0.0.0] 10.32.148.0
     Read/Write? (true, t, false, f): [true] f
     Access host subnet area in dot notation: [0.0.0.0]
     Read/Write? (true, t, false, f): [true]
     Access host subnet area in dot notation: [0.0.0.0] 10.33.0.0
     Read/Write? (true, t, false, f): [true] f
     Access host subnet area in dot notation: [0.0.0.0]
     Read/Write? (true, t, false, f): [true]
     Access host subnet area in dot notation: [0.0.0.0]
     Read/Write? (true, t, false, f): [true]
     Committing configuration...done.
```

Sample mibCapability configuration

```
switch:admin> snmpconfig --show mibCapability
     FA-MIB: YES
     FICON-MIB: YES
     HA-MIB: YES
     SW-TRAP: YES
        swFCPortScn: YES
        swEventTrap: YES
        swFabricWatchTrap: YES
        swTrackChangesTrap: NO
     FA-TRAP: YES
        connUnitStatusChange: YES
        connUnitEventTrap: NO
        connUnitSensorStatusChange: YES
        connUnitPortStatusChange: YES
     SW-EXTTRAP: NO
     FICON-TRAP: NO
     HA-TRAP: YES
        fruStatusChanged: YES
        cpStatusChanged: YES
        fruHistoryTrap: NO
```

Sample systemGroup configuration (default)

```
switch:admin> snmpconfig --default systemGroup
     *****
     This command will reset the agent's system group configuration back to factory
default
     *****
               sysDescr = Fibre Channel Switch
            sysLocation = End User Premise
             sysContact = Field Support
              authTraps = 0 (OFF)

     *****
     Are you sure? (yes, y, no, n): [no] y
```

Using Legacy Commands for SNMPv1

You should use the **snmpConfig** command to configure the SNMPv1 agent and traps (see "Using the snmpConfig Command" on page 4-16). However, if necessary for backward compatibility, you can choose to use legacy commands.

Sample SNMP agent configuration information

```
switch:admin> agtcfgshow
Current SNMP Agent Configuration
         Customizable MIB-II system variables:
                 sysDescr = FC Switch
              sysLocation = End User Premise
               sysContact = Field Support.
                authTraps = 1 (ON)

     SNMPv1 community and trap recipient configuration:
       Community 1: Secret C0de (rw)
         Trap recipient: 192.168.1.51
         Trap recipient Severity level: 4
       Community 2: OrigEquipMfr (rw)
         Trap recipient: 192.168.1.26
         Trap recipient Severity level: 0
       Community 3: private (rw)
         No trap recipient configured yet
       Community 4: public (ro)
         No trap recipient configured yet
       Community 5: common (ro)
         No trap recipient configured yet
       Community 6: FibreChannel (ro)
         No trap recipient configured yet

     SNMP access list configuration:
     Entry 0:  Access host subnet area 192.168.64.0 (rw)]
     Entry 1:  No access host configured yet
     Entry 2:  No access host configured yet
     Entry 3:  No access host configured yet
     Entry 4:  No access host configured yet
     Entry 5:  No access host configured yet
```

Sample modification of the SNMP configuration values

```
switch:admin> agtcfgset

     Customizing MIB-II system variables ...

     At each prompt, do one of the followings:
       o <Return> to accept current value,
       o enter the appropriate new value,
       o <Control-D> to skip the rest of configuration, or
       o <Control-C> to cancel any change.

     To correct any input mistake:
     <Backspace> erases the previous character,
     <Control-U> erases the whole line,
     sysDescr: [FC Switch]
     sysLocation: [End User Premise]
     sysContact: [Field Support.]
     authTrapsEnabled (true, t, false, f): [true]

     SNMP community and trap recipient configuration:
     Community (rw): [Secret C0de]
     Trap Recipient's IP address in dot notation: [192.168.1.51]
     Trap recipient Severity level : (0..5) [0] 3
     Community (rw): [OrigEquipMfr]
     Trap Recipient's IP address in dot notation: [192.168.1.26]
     Trap recipient Severity level : (0..5) [0]
     Community (rw): [private]
     Trap Recipient's IP address in dot notation: [0.0.0.0] 192.168.64.88
     Trap recipient Severity level : (0..5) [0] 1
     Community (ro): [public]
     Trap Recipient's IP address in dot notation: [0.0.0.0]
     Community (ro): [common]
     Trap Recipient's IP address in dot notation: [0.0.0.0]
     Community (ro): [FibreChannel]
     Trap Recipient's IP address in dot notation: [0.0.0.0]

     SNMP access list configuration:
     Access host subnet area in dot notation: [0.0.0.0] 192.168.64.0
     Read/Write? (true, t, false, f): [true]
     Access host subnet area in dot notation: [0.0.0.0]
     Read/Write? (true, t, false, f): [true]
     Access host subnet area in dot notation: [0.0.0.0]
     Read/Write? (true, t, false, f): [true]
     Access host subnet area in dot notation: [0.0.0.0]
     Read/Write? (true, t, false, f): [true]
     Access host subnet area in dot notation: [0.0.0.0]
     Read/Write? (true, t, false, f): [true]
     Access host subnet area in dot notation: [0.0.0.0]
     Read/Write? (true, t, false, f): [true]
     Committing configuration...done.
     value = 1 = 0x1
```

Sample reset of the SNMP agent configuration to default values

```
switch:admin> agtcfgdefault
*****
This command will reset the agent's configuration back to factory default
*****
     Current SNMP Agent Configuration
     Customizable MIB-II system variables:
               sysDescr = Fibre Channel Switch.
            sysLocation = End User Premise
              sysContact = sweng
                authTraps = 0 (OFF)
SNMPv1 community and trap recipient configuration:
       Community 1: Secret C0de (rw)
         Trap recipient: 192.168.15.41
         Trap recipient Severity level: 4
       Community 2: OrigEquipMfr (rw)
         No trap recipient configured yet
       Community 3: private (rw)
         No trap recipient configured yet
       Community 4: public (ro)
         No trap recipient configured yet
       Community 5: common (ro)
         No trap recipient configured yet
       Community 6: FibreChannel (ro)
         No trap recipient configured yet
SNMP access list configuration:
     Entry 0:  Access host subnet area 192.168.64.0 (rw)]
     Entry 1:  No access host configured yet
     Entry 2:  No access host configured yet
     Entry 3:  No access host configured yet
     Entry 4:  No access host configured yet
     Entry 5:  No access host configured yet
*****
     Are you sure? (yes, y, no, n): [no] y
     Committing configuration...done.
     agent configuration reset to factory default
     Current SNMP Agent Configuration
     Customizable MIB-II system variables:
               sysDescr = Fibre Channel Switch.
             sysLocation = End User Premise
              sysContact = Field Support.
                authTraps = 0 (OFF)
SNMPv1 community and trap recipient configuration:
       Community 1: Secret C0de (rw)
         No trap recipient configured yet
       Community 2: OrigEquipMfr (rw)
         No trap recipient configured yet
       Community 3: private (rw)
         No trap recipient configured yet
       Community 4: public (ro)
         No trap recipient configured yet
       Community 5: common (ro)
         No trap recipient configured yet
       Community 6: FibreChannel (ro)
         No trap recipient configured yet
(output truncated)
```

Sample modification of the options for configuring SNMP MIB traps

```
switch:admin> snmpmibcapset
The SNMP Mib/Trap Capability has been set to support
FE-MIB
SW-MIB
FA-MIB
FA-TRAP
FA-MIB (yes, y, no, n): [yes]
FICON-MIB (yes, y, no, n): [no] y
HA-MIB (yes, y, no, n): [no] y
SW-TRAP (yes, y, no, n): [no] y
   swFCPortScn (yes, y, no, n): [no]
   swEventTrap (yes, y, no, n): [no]
   swFabricWatchTrap (yes, y, no, n): [no]
   swTrackChangesTrap (yes, y, no, n): [no]
FA-TRAP (yes, y, no, n): [yes]
   connUnitStatusChange (yes, y, no, n): [no]
   connUnitEventTrap (yes, y, no, n): [no]
   connUnitSensorStatusChange (yes, y, no, n): [no]
   connUnitPortStatusChange (yes, y, no, n): [no]
SW-EXTTRAP (yes, y, no, n): [no] y
FICON-TRAP (yes, y, no, n): [no] y
   linkRNIDDeviceRegistration (yes, y, no, n): [no]
   linkRNIDDeviceDeRegistration (yes, y, no, n): [no]
   linkLIRRListenerAdded (yes, y, no, n): [no]
   linkLIRRListenerRemoved (yes, y, no, n): [no]
   linkRLIRFailureIncident (yes, y, no, n): [no]
HA-TRAP (yes, y, no, n): [no] y
   fruStatusChanged (yes, y, no, n): [no]
   cpStatusChanged (yes, y, no, n): [no]
   fruHistoryTrap (yes, y, no, n): [no]
Avoid-Duplicate-TRAP (yes, y, no, n): [no] y
switch:admin>
```

The following **snmpMibCapSet** parameters for FA-TRAP appear in the preceding example:

- connUnitStatusChange indicates that the overall status of the connectivity unit has changed. Its variables are:
  - connUnitStatus: the status of the connection unit
  - connUnitState: the state of the connection unit
- connUnitEventTrap indicates that the connectivity unit has generated an event. Its variables are:
  - connUnitEventId: the internal event ID
  - connUnitEventType: the type of this event
- connUnitEventObject is used with the connUnitEventType to identify the object to which the event refers.
- connUnitEventDescr is the description of the event.
- connUnitSensorStatusChange indicates that the status of the sensor associated with the connectivity unit has changed.
- connUnitSensorStatus is the status indicated by the sensor.
- connUnitPortStatusChange indicates that the status of the sensor associated with the connectivity unit has changed.
- connUnitPortStatus shows overall protocol status for the port.
- connUnitPortState shows the user-specified state of the port hardware.

Sample view of the SNMP MIB trap setup

```
switch:admin> snmpmibcapshow
  FA-MIB: YES
  FICON-MIB: YES
  HA-MIB: YES
  SW-TRAP: YES
     swFCPortScn: YES
     swEventTrap: YES
     swFabricWatchTrap: YES
     swTrackChangesTrap: YES
  FA-TRAP: YES
  SW-EXTTRAP: YES
  HA-TRAP: YES
     fruStatusChanged: YES
     cpStatusChanged: YES
     fruHistoryTrap: YES
```

Configuring Secure File Copy

You can use the **configure** command to specify that secure file copy (scp) be used for configuration uploads and downloads.

For example:

```
switch:admin> configure

Not all options will be available on an enabled switch.
To disable the switch, use the "switchDisable" command.

Configure...

  System services (yes, y, no, n): [no] n
  ssl attributes (yes, y, no, n): [no] n
  http attributes (yes, y, no, n): [no] n
  snmp attributes (yes, y, no, n): [no] n
  rpcd attributes (yes, y, no, n): [no] n
  cfgload attributes (yes, y, no, n): [no] y

        Enforce secure config Upload/Download (yes, y, no, n): [no] y
switch:admin>
```

Chapter

5

Configuring Advanced Security

This chapter provides information and procedures for configuring advanced Fabric OS v5.2.0 security feature, Access Control Lists (ACL) policies for FC port and switch binding.

Note

Run all commands in this chapter by logging in to Administrative Domain (AD) 255 or if Administrative Domains have not been implemented log in to AD 0.

This chapter covers the following topics:

- "About Access Control List (ACL) Policies" on page 5-2
- "Configuring ACL Policies" on page 5-3
- "Distributing the Policy Database" on page 5-11

For information about licensed security features available in Secure Fabric OS, see the *Secure Fabric OS Administrator's Guide*.

About Access Control List (ACL) Policies

Fabric OS provides the following policies:

- **Device Connection Control** (DCC) policies—Used to restrict which Fibre Channel device ports can connect to which Fibre Channel switch ports.
- **Switch Connection Control** (SCC) policy—Used to restrict which switches can join the switch.

Each supported policy is identified by a specific name, and only one policy of each type can exist (except for DCC policies). Policy names are case sensitive and must be entered in all uppercase.

How the ACL Policies are Stored

The policy are stored in a local database. The database contains both ACL policies types (SCC and DCC). The policy are grouped by state and type.

A policy can be in the following state:

- **Active**—The policy is being enforced by the switch.
- **Defined**—The policy has been set up but is not enforced.

A group of policies is called a Policy Set.

Each switch has the following two sets:

- **Active policy set—**Contains ACL policies being enforced by the switch.
- **Defined policy set—**Contains a copy of all ACL policies on the switch.

When you activate a policy, the defined policy either replaces the policy with the same name in the active set or becomes a new active policy. If a policy appears in the defined set but not in the active set, the policy was saved but has not been activated. If a policy with the same name appears in both the defined and active sets but they have different values, then the policy has been modified but the changes have not been activated.

Identifying Policy Members

Specify policy members by device port WWN, switch WWN, domain IDs, or switch names, depending on the policy. The valid methods for specifying policy members are listed in Table 5-1.

Table 5-1 Valid Methods for Specifying Policy Members

Policy Name	Device Port WWN	Switch WWN	Domain ID	Switch Name
DCC_POLICY_*nnn*	Yes	Yes	Yes	Yes
SCC_POLICY	No	Yes	Yes	Yes

Configuring ACL Policies

All policy modifications are saved in volatile memory until those changes are saved or activated. You can create multiple sessions to the switch from one or more hosts. However, Fabric OS allows only one ACL transaction at a time. If a second ACL transaction is started, it fails.

The Secure Fabric OS and Fabric OS SCC and DCC policies are not interchangeable. Uploading and saving a copy of the Fabric OS configuration after creating policies is strongly recommended. Use the **configUpload** command to upload a copy of the configuration file. For more information about this command, see the *Fabric OS Command Reference*.

Note

All changes, including the creation of new policies, are saved and activated on the local switch only—unless the switch is in a fabric that has a strict or tolerant fabric-wide consistency policy for the ACL policy type (SCC or DCC). See "Distributing the Policy Database" on page 5-11 for more information on the database settings and fabric-wide consistency policy.

Use the instructions in the following sections to manage DCC and SCC policies:

- "Displaying ACL Policies" on page 5-4
- Displays a list of all active and defined ACL policies on the switch.
- "Configuring a DCC Policy" on page 5-4
  Multiple DCC policies can be created using the naming convention DCC_POLICY_*nnn*, with *nnn* representing a unique string.
- "Creating an SCC Policy" on page 5-7
  One SCC policy can be created.
- "Saving Changes to ACL Policies" on page 5-8
  Save changes to memory without actually implementing the changes within the fabric or to the switch. This saved but inactive information is known as the "defined policy set."
- "Activating Changes to ACL Policies" on page 5-8
  Simultaneously save and implement all the policy changes made since the last time changes were activated. The activated policies are known as the "active policy set."
- "Adding a Member to an Existing Policy" on page 5-9
  Add one or more members to a policy. The aspect of the fabric covered by each policy is closed to access by all devices/switches that are not listed in that policy.
- "Removing a Member from a Policy" on page 5-9
  Remove one or more members from a policy. If all members are removed from a policy, that aspect of the fabric becomes closed to all access.
- "Deleting a Policy" on page 5-10
  Delete an entire policy; deleting a policy opens up that aspect of the fabric to all access.

Displaying ACL Policies

Use the secPolicyShow command to display the Active and Defined policy sets. The following example shows a switch that has no SCC and DCC policies.

secPolicyShow displays the following information:

- **Active Policy Set**—The policies that are being enforced.
- **Defined Policy Set**—The policies that have been saved. Policies created in the same login session also appear but these policies are automatically deleted if the user logs out without saving.

To display the ACL policies

1. Connect to the switch and log in.
2. Type the **secPolicyShow** command:

```
switch:admin> secPolicyShow


_________________________________________________
                ACTIVE POLICY SET




_________________________________________________
                DEFINED POLICY SET

switch:admin>
```

Note

The defined policy set list is described in "Saving Changes to ACL Policies" on page 5-8 and active policy set is described in "Activating Changes to ACL Policies" on page 5-8.

Configuring a DCC Policy

Multiple DCC policies can be used to restrict which device ports can connect to which switch ports. The devices can be initiators, targets, or intermediate devices such as SCSI routers and loop hubs. By default, all device ports are allowed to connect to all switch ports; no DCC policies exist until they are created.

Each device port can be bound to one or more switch ports; the same device ports and switch ports might be listed in multiple DCC policies. After a switch port is specified in a DCC policy, it permits connections only from designated device ports. Device ports that are not specified in any DCC policies are allowed to connect only to switch ports that are not specified in any DCC policies.

When a DCC violation occurs, the related port is automatically disabled and must be re-enabled using the **portEnable** command.

The procedure used to create a DCC policy is described after Table 5-2, which shows the possible DCC policy states.

Table 5-2 DCC Policy States

Policy State	Characteristics
No policy	Any device can connect to any switch port in the fabric.
Policy with no entries	Any device can connect to any switch port in the fabric. An empty policy is the same as no policy.
Policy with entries	If a device WWN is specified in a DCC policy, that device is only allowed access to the switch if connected by a switch port listed in the same policy. If a switch port is specified in a DCC policy, it only permits connections from devices that are listed in the policy. Devices with WWNs that are not specified in a DCC policy are allowed to connect to the switch at any switch ports that are not specified in a DCC policy. Switch ports and device WWNs may exist in multiple DCC policies. Proxy devices are always granted full access and can connect to any switch port in the fabric.

DCC Policy Restrictions

The following restrictions apply when using DCC policies:

- Fabric OS v5.2.0 supports DCC policies. You cannot directly transfer DCC policies created in Secure Fabric OS to policies to be used in Fabric OS.

  Policies created in Secure Fabric OS are deleted when Secure Fabric OS is disabled; policies created in Fabric OS are deleted when Secure Fabric OS is enabled. Therefore, back up DCC policies before enabling or disabling Secure Fabric OS.
- Some older private-loop HBAs do not respond to port login from the switch and are not enforced by the DCC policy. This does not create a security problem because these HBAs cannot contact any device outside of their immediate loop.
- DCC policies cannot manage or restrict iSCSI connections, that is an FC Initiator connection from an iSCSI gateway.
- You cannot manage proxy devices with DCC policies. Proxy devices are always granted full access, even if the DCC policy has an entry that restricts or limits access of a proxy device.

Creating a DCC Policy

DCC policies must follow the naming convention “DCC_POLICY_*nnn*,” where *nnn* represents a unique string. To save memory and improve performance, one DCC policy per switch or group of switches is recommended.

Device ports must be specified by port WWN. Switch ports can be identified by the switch WWN, domain ID, or switch name followed by the port or area number. To specify an allowed connection, enter the device port WWN, a semicolon, and the switch port identification.

Following are the possible methods of specifying an allowed connection:

- *deviceportWWN*;*switchWWN* (port or area number)
- *deviceportWWN*;*domainID* (port or area number)
- *deviceportWWN*;*switchname* (port or area number)

To create a DCC policy

1. Connect to the switch and log in.
2. Type **secPolicyCreate "DCC_POLICY_***nnn***", "***member***;...;***member***"**.

   DCC_POLICY_*nnn* is the name of the DCC policy; *nnn* is a string consisting of up to 19 alphanumeric or underscore characters to differentiate it from any other DCC policies.

   The *member* contains device or switch port information: *deviceportWWN*;*switch(port) where:*

deviceportWWN	WWN of the device port.
switch	Either the switch WWN, domain ID, or switch name. The port can be specified by port or area number. Designating ports automatically includes the devices currently attached to those ports. The ports can be specified using any of the following syntax methods: (*) Selects all ports on the switch. (1-6) Selects ports 1 through 6. [*] Selects all ports and all devices attached to those ports. [3, 9] Selects ports 3 and 9 and all devices attached to those ports. [1-3, 9] Selects ports 1, 2, 3, 9, and all devices attached to those ports.

3. To save or activate the new policy, enter either the **secPolicySave** or the **secPolicyActivate** command.

   If neither of these commands is entered, the changes are lost when the session is logged out. For more information about these commands, see "Saving Changes to ACL Policies" on page 5-8 and "Activating Changes to ACL Policies" on page 5-8.

Examples of Creating DCC Policies

To create the DCC policy "DCC_POLICY_server" that includes device 11:22:33:44:55:66:77:aa and port 1 and port 3 of switch domain 1:

```
switch:admin> secpolicycreate "DCC_POLICY_server", "11:22:33:44:55:66:77:aa;1(1,3)"
DCC_POLICY_server has been created
```

To create the DCC policy “DCC_POLICY_storage” that includes device port WWN 22:33:44:55:66:77:11:bb, all ports of switch domain 2, and all currently connected devices of switch domain 2:

```
switch:admin> secpolicycreate "DCC_POLICY_storage", "22:33:44:55:66:77:11:bb;2[*]"
DCC_POLICY_storage has been created
```

To create the DCC policy “DCC_POLICY_abc” that includes device 33:44:55:66:77:11:22:cc and ports 1 through 6 and port 9 of switch domain 3:

```
switch:admin> secpolicycreate "DCC_POLICY_abc", "33:44:55:66:77:11:22:cc;3(1-6,9)"
DCC_POLICY_xxx has been created
```

To create the DCC policy “DCC_POLICY_example” that includes devices 44:55:66:77:22:33:44:dd and 33:44:55:66:77:11:22:cc, ports 1 through 4 of switch domain 4, and all devices currently connected to ports 1 through 4 of switch domain 4:

```
switch:admin> secpolicycreate "DCC_POLICY_example",
"44:55:66:77:22:33:44:dd;33:44:55:66:77:11:22:cc;4[1-4]"
DCC_POLICY_xxx has been created
```

Creating an SCC Policy

Fabric OS v5.2.0 and higher supports a SCC policy in Fabric OS. An SCC policy created in Secure Fabric OS cannot directly transfered to Fabric OS. Policies created in Secure Fabric OS are deleted when Secure Fabric OS is disabled; policies created in Fabric OS are deleted when Secure Fabric OS is enabled. Back up SCC policies before enabling or disabling Secure Fabric OS.

The SCC policy is used to restrict which switches can join the fabric. Switches are checked against the policy each time an E_Port-to-E_Port connection is made. The policy is named SCC_POLICY and accepts members listed as WWNs, domain IDs, or switch names. Only one SCC policy can be created.

By default, any switch is allowed to join the fabric; the SCC policy does not exist until it is created. When connecting a Fibre Channel router to a fabric or switch that has an active SCC policy, the front domain of the Fibre Channel router must be included in the SCC policy.

SCC policy states are shown in Table 5-3.

Table 5-3 SCC Policy States

Policy State	SCC Policy Enforcement
No active policy	All switches can connect to the switch with the specified policy.
Active policy that has no members	All neighboring switches are segmented.
Active policy that has members	The neighboring switches not specified in the SCC policy are segmented.

To create an SCC policy

1. Connect to the switch and log in.
2. Type **secPolicyCreate "SCC_POLICY", "***member***;...;***member***"**.

   member indicates a switch that is permitted to join the fabric. Specify switches by WWN, domain ID, or switch name. Enter an asterisk (*) to indicate all the switches in the fabric.

   For example, to create an SCC policy that allows switches that have domain IDs 2 and 4 to join the fabric:

```
switch:admin> secpolicycreate "SCC_POLICY", "2;4"
SCC_POLICY has been created
```

3. To save or activate the new policy, enter either the **secPolicySave** or the **secPolicyActivate** command.

   If neither of these commands is entered, the changes are lost when the session is logged out. For more information about these commands, see "Saving Changes to ACL Policies" on page 5-8 and "Activating Changes to ACL Policies" on page 5-8.

Saving Changes to ACL Policies

You can save changes to ACL policies without activating them by entering the **secPolicySave** command. This saves the changes to the defined policy set. Until the **secPolicySave** or **secPolicyActivate** command is issued, all policy changes are in volatile memory only and are lost if the switch reboots or the current session is logged out.

To save changes without activating the policies

1. Connect to the switch and log in.
2. Type the **secPolicySave** command.

```
switch:admin> secpolicysave
```

Activating Changes to ACL Policies

Implement changes to the ACL policies using the **secPolicyActivate** command. This saves the changes to the active policy set and activates all policy changes since the last time the command was issued. You cannot activate policies on an individual basis; all changes to the entire policy set are activated by the command. Until a **secPolicySave** or **secPolicyActivate** command is issued, all policy changes are in volatile memory only and are lost upon rebooting.

To activate changes

1. Connect to the switch and log in.
2. Type the **secPolicyActivate** command:

```
switch:admin> secpolicyactivate
About to overwrite the current Active data.
ARE YOU SURE (yes, y, no, n): [no] y
```

Adding a Member to an Existing Policy

Add members to the ACL policies by using the **secPolicyAdd** command. As soon as a policy has been activated, the aspect of the fabric managed by that policy is enforced.

To add a member to an existing ACL policy

1. Connect to the switch and log in.
2. Type **secPolicyAdd "***policy_name***", "***member***;...;***member***"**.

   where *policy_name* is the name of the ACL policy, *member* is the item to be added to the policy as identified by IP address, switch domain ID, device or switch WWN, or switch name.
3. To implement the change immediately, enter the **secPolicyActivate** command.

   For example, to add a member to the SCC_POLICY using the switch WWN:

```
switch:admin> secpolicyadd "SCC_POLICY", "12:24:45:10:0a:67:00:40"
Member(s) have been added to SCC_POLICY.
```

   To add two devices to the DCC policy, and to attach domain 3 ports 1 and 3 (WWNs of devices are 11:22:33:44:55:66:77:aa and 11:22:33:44:55:66:77:bb):

```
switch:admin> secpolicyadd "DCC_POLICY_abc",
"11:22:33:44:55:66:77:aa;11:22:33:44:55:66:77:bb;3(1,3)"
```

Removing a Member from a Policy

To remove a member from an ACL policy

1. Connect to the switch and log in.
2. Type **secPolicyRemove "***policy_name***", "***member***;...;***member***"**.

   where *policy_name* is the name of the ACL policy. *member* is the device or switch to be removed from the policy, identified by IP address, switch domain ID, device or switch WWN, or switch name.
3. To implement the change immediately, enter the **secPolicyActivate** command.

   For example, to remove a member that has a WWN of 12:24:45:10:0a:67:00:40 from SCC_POLICY:

```
switch:admin> secpolicyremove "SCC_POLICY", "12:24:45:10:0a:67:00:40"
Member(s) have been removed from SCC_POLICY.
```

Deleting a Policy

To delete an ACL policy

1. Connect to the switch and log in.
2. Type **secPolicyDelete "***policy_name***"**.

   where *policy_name* is the name of the ACL policy.
3. To implement the change immediately, enter the **secPolicyActivate** command.

```
switch:admin> secpolicydelete "DCC_POLICY_ALL"
About to delete policy Finance_Policy.
Are you sure (yes, y, no, n):[no] y
Finance_Policy has been deleted.
```

Aborting All Uncommitted Changes

Use the **secPolicyAbort** command to abort all ACL policy changes that have not yet been saved.

To abort all unsaved changes

1. Connect to the switch and log in.
2. Type the **secPolicyAbort** command:

```
switch:admin> secpolicyabort
Unsaved data has been aborted.
```

   All changes since the last time the **secPolicySave** or **secPolicyActivate** commands were entered are aborted.

Distributing the Policy Database

Fabric OS lets you manage and enforce the ACL policy database on either a per-switch or fabric-wide basis. The local switch distribution setting and the fabric-wide consistency policy affect the switch's ACL policy database and related distribution behavior.

The ACL policy database is managed as follows:

- **Switch database distribution setting—**Controls whether or not the switch accepts or rejects databases distributed from other switches in the fabric. The distribution command sends the database from one switch to another, overwriting the target switches database with the distributed one. To send or receive a database the setting must be accept. For configuration instructions, see "Configuring the Database Distribution Settings" on page 5-12.
- **Manually distribute an ACL policy database**—Run the distribute command to push the local database of the specified policy type to target switches. "Distributing ACL Policies to other Switches" on page 5-13
- **Fabric-wide consistency policy**—Use to ensure that switches in the fabric enforce the same policies. Set a strict or tolerant fabric-wide consistency policy for each ACL policy type to automatically distribute that database when a policy change is activated. If a fabric-wide consistency policy is not set, then the policies are managed on per switch basis. For configuration instructions, see"Setting the Consistency Policy Fabric-Wide" on page 5-14

The following table explains the how the local database distribution settings and the fabric-wide consistency policy affect the local database when the switch is the target of a distribution command.

Table 5-4 Interaction between Fabric-wide Consistency Policy and Distribution Settings

Distribution Setting	Fabric-wide Consistency Policy		
	Absent (default)	**Tolerant**	**Strict**
Reject	Database is protected, it cannot be overwritten. Might not match other databases in the fabric.	Invalid configuration.[a]	Invalid configuration.[a]
Accept (default)	Database is not protected, the database can be overwritten. If the switch initiating a distribute command has a strict or tolerant fabric-wide consistency policy, the fabric-wide policy is also overwritten. Might not match other databases in the fabric.	Database is not protected. Automatically distributes activated changes to other v5.20 switches in fabric. Allows switches running Fabric OS v5.1.x and earlier in fabric. Might not match other databases in the fabric.	Database is not protected. Automatically distributes activated changes to all switches in the fabric. Fabric can only contain switches running Fabric OS v5.2.0 or higher. Active database is the same for all switches in fabric.

a. Error returned indicating that the distribution setting must be accept before you can set the fabric-wide consistency policy.

Configuring the Database Distribution Settings

The distribution settings control whether a switch accepts or rejects distributions of databases from other switches and whether or not the switch may initiate a distribution. Configure the distribution setting to reject when maintaining the database on a per-switch basis.

The following table lists the databases supported starting in Fabric OS v 5.2.0.

Table 5-5 Supported Databases Starting in Fabric OS v5.2.0

Database Type	Database Identifier (ID)
SCC policy database	SCC
DCC policy database	DCC
Password database	PWD

Note

See "Distributing the Local User Database" on page 3-12 for details on managing and distributing the user account database.

To display the database distribution settings

1. Connect to the switch.
2. Enter the following command:

```
switch:admin> fddcfg --showall
Local Switch Configuration for all Databases:-
DATABASE - Accept/Reject
-------------------------
SCC      - accept
DCC      - accept
PWD      - accept
Fabric Wide Consistency Policy:- ""
```

To enable switch-local protection

1. Connect to the switch.
2. Enter the following command:

```
fddCfg --localreject <database_ID>
```

localreject	Refuse the databases distributed from other switches. Cannot distribute local database, manually or automatically, to other switches.
database_id	A semicolon-separated list of the local databases to be distributed, either SCC and/or DCC.

To disable switch-local protection

1. Connect to the switch.
2. Enter the following command:

```
fddCfg --localaccept <database_ID>
```

localaccept	Default setting. Allows local database to be overwritten with databases received from other switches. Allows local database to be manually or automatically distributed to other switches.
database_id	A semicolon-separated list of the local databases to be distributed, either SCC and/or DCC.

Distributing ACL Policies to other Switches

This section explains how to manually distribute local ACL policy databases to other Fabric v5.2.0 and higher switches. The distribute command has the following dependencies:

- All target switches must be running Fabric OS v5.2.0 or higher.
- All target switches must accept the database distribution (see "Configuring the Database Distribution Settings" on page 5-12).
- The fabric must have a tolerant or no (absent) fabric-wide consistency policy (see "Setting the Consistency Policy Fabric-Wide" on page 5-14).

  If the fabric-wide consistency policy for a database is strict, the database cannot be manually distributed. When you set a strict fabric-wide consistency policy for a database, the distribution mechanism is automatically invoked whenever the database changes.
- The local distribution setting must be accepted. To be able to initiate the distribute command, set the local distribution to accept.

Table 5-6 describes how the target switch database distribution settings affect the distribution:

Table 5-6 ACL Policy Database Distribution Behavior

Target Switch		Distribution	Results
Fabric OS Version	**Database setting**		
v5.1.0 or earlier	NA	Fails	An error is returned. The entire transaction is aborted and no databases are updated.
v5.2.0	Reject	Fails	The target switch explicitly refuses the distribution. The entire transaction is aborted and no databases are updated.
	Accept	Succeeds	The target switch accepts the distribution.

To distribute the local ACL policies

1. Connect to the switch.
2. Enter the following command:

```
distribute -p <database_id> -d <switch_list>
```

database_id	A semicolon-separated list of the local databases to be distributed: SCC and/or DCC.
switch_list	A is a semicolon-separated list of switch Domain IDs, switch names, or switch WWN addresses of the target switches that will received the distribution. Use an asterisk (*) to distribution the database to all Fabric OS v5.2.0 and higher switches in the fabric. For example, entering the command: **distribute -p SCC -d "*"** distributes the SCC policy to all v5.2.0 and higher switches in the fabric.

Setting the Consistency Policy Fabric-Wide

Fabric OS v5.2.0 introduces a fabric-wide consistency policy for SCC and DCC policies. The fabric-wide consistency policy enforcement setting determines the distribution behavior when changes to a policy are activated. Using the tolerant or strict fabric-wide consistency policy ensures that changes to local ACL policy databases are automatically distributed to other switches in the fabric.

When you set the fabric-wide consistency policy using the **`fddCfg`** command with the **`--fabwideset <database_id>`** option, both the fabric-wide consistency policy and specified database are distributed to the fabric.The active policies of the specified databases overwrite the corresponding active and defined policies on the target switches.

Policy changes that are saved but not activated are stored locally until a policy database change is activated. Activating a policy automatically distributes the Active policy set for that policy type (SCC and/or DCC), to the other switches in the fabric.

Note

FC routers cannot join a fabric with a strict fabric-wide consistency policy. FC routers do not support the fabric-wide consistency policies.

Table 5-7 describes the fabric-wide consistency settings.

Table 5-7 Fabric-wide Consistency Policy Settings

Setting	Value	When a policy is activated
Absent	null	Database is not automatically distributed to other switches in the fabric.
Tolerant	*database_id*	All updated and/or new policies of the type specified (SCC and/or DCC) are distributed to all Fabric v5.2.0 and higher switches in the fabric.
Strict	*database_id*:S	All updated and/or new policies of the type specified (SCC and/or DCC) are distributed to all switches in the fabric.

To display the fabric-wide consistency policy

1. Connect to the switch and log in.
2. Enter the **fddCfg --showall** command.

   The following example shows a not defined fabric-wide consistency policy.

```
switch:admin> fddcfg --showall
Local Switch Configuration for all Databases:-
DATABASE - Accept/Reject
-------------------------
SCC      - accept
DCC      - accept
PWD      - accept
Fabric Wide Consistency Policy:- ""
```

To set the fabric-wide consistency policy

1. Connect to the switch and log in.
2. Enter the following command:

```
fddCfg --fabwideset "<policy_ID>"
```

   Where *policy_ID* is a semicolon-separated list *database_setting*;*database_setting* equal to

null	Exclude the database ID from the list to set the policy to absent.
database_id	Sets a tolerant policy for a database. The database ID is either SCC or DCC.
database_id:S	Sets the policy to strict. The database ID is either SCC or DCC.

   The following examples show how to set a strict SCC and tolerant DCC fabric-wide consistency policy.

```
switch:admin> fddcfg --fabwideset "SCC:S;DCC"
```

```
switch:admin> fddcfg --showall
Local Switch Configuration for all Databases:-
DATABASE - Accept/Reject
-------------------------
SCC      - accept
DCC      - accept
PWD      - accept
Fabric Wide Consistency Policy:- "SCC:S;DCC"
```

Notes on Joining a Switch to the Fabric

When a switch is joined to a fabric with a tolerant SCC or DCC fabric-wide consistency policy, the joining switch must have a matching tolerant SCC or DCC fabric-wide consistency policy. If the tolerant SCC or DCC fabric-wide consistency policies do not match, the switch can join the fabric, but an error message flags the mismatch. If the tolerant SCC and DCC fabric-wide consistency policies match, the corresponding SCC and DCC ACL policies are compared.

The enforcement of fabric-wide consistency policy involves comparison of only the Active policy set. If the ACL policies match, the switch joins the fabric successfully. If the ACL policies are absent on the switch or on the fabric, the switch joins the fabric successfully, and the ACL policies are copied automatically from where they are present side to where they are absent. The Active policies set where they are present overwrite the Active and Defined policies set where they are absent. If the ACL policies do not match, the switch can join the fabric, but an error message flags the mismatch.

Under both conflicting conditions, **secPolicyActivate** is blocked in the merged fabric.Uuse **fddcfg –fabwideset** command to resolve the fabric-wide consistency policy conflicts. Use the **distribute** command to explicitly resolve conflicting ACL policies.

When a switch is joined to a fabric with a strict SCC or DCC fabric-wide consistency policy, the joining switch must have a matching fabric-wide consistency policy. If the strict SCC or DCC fabric-wide consistency policies do not match, the switch cannot join the fabric and the neighboring E_ports will be disabled. If the strict SCC and DCC fabric-wide consistency policies match, the corresponding SCC and DCC ACL policies are compared.

The enforcement of fabric-wide consistency policy involves comparison of only the Active policy set.If the ACL polices match, the switch joins the fabric successfully. If the ACL policies are absent either on the switch or on the fabric, the switch joins the fabric successfully, and the ACL policies are copied automatically from where they are present to where they are absent. The Active policy set where it is present overwrites the Active and Defined policy set where it is absent. If the ACL policies do not match, the switch cannot join the fabric and the neighboring E_ports will be disabled.

Use the **fddcfg –fabwideset** command on either this switch or the fabric to set a matching strict SCC or DCC fabric-wide consistency policy. Use ACL policy commands to delete the conflicting ACL policy from one side to resolve ACL policy conflict. If neither the fabric nor the joining switch is configured with a fabric-wide consistency policy, there are no ACL merge checks required.

The descriptions above also applies to joining two fabrics. In this context, the joining switch in the above descriptions becomes a joining fabric.

Note

The existing SCC policy in the adjacent or joining switch are enforced before the fabric-wide consistency policy is applied.

The next sections describe the interaction between the databases with active SCC and DCC policies and combinations of fabric-wide consistency policy settings when fabrics are merged.

Matching fabric-wide consistency policies

For example: Fabric A with SCC:S;DCC (strict SCC and tolerant DCC) joins Fabric B with SCC:S;DCC (strict SCC and tolerant DCC), the fabrics can merge as long as the SCC policies match (both are strict).

Table 5-8 describes the impact of merging fabrics with the same fabric-wide consistency policy that have SCC and/or DCC policies.

Table 5-8 Merging Fabrics with Matching Fabric-wide Consistency Policies

Fabric-wide Consistency Policy	Fabric A ACL Policies	Fabric B ACL Policies	Merge Results	Database Copied
None	None	None	Succeeds	No ACL policies copied.
	None	SCC/DCC	Succeeds	No ACL policies copied.
Tolerant	None	None	Succeeds	No ACL policies copied.
	None	SCC/DCC	Succeeds	ACL policies are copied from B to A.
	SCC/DCC	SCC/DCC	Succeeds	If A and B policies don't match, a warning displays and policy commands are disabled[a].
Strict	None	None	Succeeds	No ACL policies copied.
	None	SCC/DCC	Succeeds	ACL policies are copied from B to A.
	Matching SCC/DCC	Matching SCC/DCC	Succeeds	No ACL policies copied.
	Different SCC/DCC policies	Different SCC/DCC policies	Fails	Ports are disabled.

a. To resolve the policy conflict, manually distribute the database you want to use to the switch with the mismatched database. Until the conflict is resolved commands such as fddcfg --fabwideset and secpolicy activate are blocked.

Non-Matching Fabric-Wide Consistency Policies

You may encounter one of the following two scenarios:

- Merging a fabric with a strict policy to a fabric with an absent, tolerant, or non-matching strict policy. The merge fails and the ports are disabled.

  The following table shows merges that are not supported:

Table 5-9 Examples of Strict Fabric Merges

Fabric-wide Consistency Policy Setting			Expected Behavior
	Fabric A	Fabric B	
Strict/Tolerant	SCC:S;DCC:S	SCC;DCC:S	Ports connecting switches are disabled.
	SCC;DCC:S	SCC:S;DCC	
Strict/Absent	SCC:S;DCC:S		
	SCC:S		
	DCC:S		
Strict/Strict	SCC:S	DCC:S	

- Merging fabrics with tolerant and absent policies. See the following table for examples:

Table 5-10 Fabric merges with Tolerant/Absent Combinations

Fabric-wide Consistency Policy Setting			Expected behavior
	Fabric A	**Fabric B**	
Tolerant/Absent	SCC;DCC		Error message logged.
	DCC		Run **fddCfg --fabwideset "<policy_ID>"** from any switch with the desired configuration to fix the conflict. The **secPolicyActivate** command is blocked until conflict is resolved.
	SCC;DCC	SCC	
	DCC	SCC	

Maintaining Configurations

Chapter 6

It is important to maintain consistent configuration settings on all switches in the same fabric, because inconsistent parameters (such as inconsistent PID formats) can cause fabric segmentation. As part of standard configuration maintenance procedures, it is recommended that you back up all important configuration data for every switch on a host computer server for emergency reference.

Note

For information about AD-enabled switches using Fabric OS v5.2.0 or higher, see Chapter 7, "Managing Administrative Domains".

This chapter contains procedures for basic switch configuration maintenance.

- "Displaying Configuration Settings" on page 6-1
- "Backing Up a Configuration" on page 6-2
- "Restoring a Configuration" on page 6-5
- "SilkWorm 48000 Configuration Form" on page 6-10

Displaying Configuration Settings

There are two ways to view configuration settings for a switch in a Brocade fabric:

- Issue the **configShow** command (less information)

  To display configuration settings using the CLI, connect to the switch, log in as admin, and enter the **configShow** command at the command line. The configuration settings vary depending on switch model and configuration.

- Upload an ASCII text file from the switch or switch module using the **configUpload** command (more information)

The uploaded switch configuration file contains the following sections and is organized as follows:

Table 6-1 Structure of an Uploaded Switch Configuration File

Section	Description
Introduction	Version of the Fabric OS from which this configuration was uploaded and the time stamp for the uploaded file creation
Boot Parameters	Variables such as the switch's name and IP address
Licenses	List of licenses that are active on the switch
Chassis Configuration Configuration	General switch configuration variables, such as diagnostic settings, fabric configuration settings, and SNMP settings

Table 6-1 Structure of an Uploaded Switch Configuration File (Continued)

Section	Description
Zoning	Zoning configuration parameters
AD-Headers	Administrative domain headers configuration parameters in files uploaded from an AD255 environment; a subsection for every AD set up (AD*n*-ZONE), where *n* is the domain number.
Defined Security Policies	List all of the defined security policies
Active Security Policies	List of all active security policies
FICU SAVED FILES	FICON CUP configuration parameters
fcRouting	Fibre Channel Routing configuration parameters
iSCSI	iSCSI configuration parameters
Banner	Security banner text

Backing Up a Configuration

In case the configuration is lost or unintentional changes are made, keep a backup copy of the configuration file (if using a SilkWorm 24000 with two domains, configuration *files*–one for each logical switch), which is what gets saved when you enter the **configUpload** command. You should keep individual backup files for all switches in the fabric. You should avoid copying configurations from one switch to another.

The following information is *not saved* in a backup:

- dnsconfig information
- passwords

If your setup supports anonymous users, and you log in as an anonymous user, password is still a required field, even though its value may be ignored by the FTP service.

You can specify the use of secure file copy (scp) during the procedure. For instructions on configuring the use of scp by default, see "Configuring Secure File Copy" on page 4-23.

Before beginning, verify that you can reach the FTP server from the switch. Using a telnet connection, save a backup copy of the configuration file from a logical switch to a host computer:

To upload a configuration file

1. Verify that the FTP service is running on the host computer.
2. Connect to the switch and log in as admin.
3. Enter the **configUpload** command. The command becomes interactive and you are prompted for the required information.

4. Respond to the prompts as follows:

Protocol (scp or ftp)	If your site requires the use of Secure Copy, specify scp. Otherwise, specify ftp.
Server Name or IP Address	Enter the name or IP address of the server where the file is to be stored; for example, 192.1.2.3. You can enter a server name if DNS is enabled. For details about the dnsConfig command, see the *Fabric OS Command Reference*.
User name	Enter the user name of your account on the server; for example, *JohnDoe*.
File name	Specify a file name for the backup file; for example, *config.txt*. Absolute path names can be specified using forward slash (/). Relative path names create the file in the user's home directory on UNIX servers, and in the directory where the FTP server is running on Windows servers.
Password	Enter your account password for the server. Note that this is a required field even if you are logged in as an anonymous user; in such cases, the value may be ignored by the FTP service.

The following example shows **configUpload** run on a switch without Admin Domains:

```
switch:admin> configupload
Protocol (scp or ftp) [ftp]: ftp
Server Name or IP Address [host]: 192.1.2.3
User Name [user]: JohnDoe
File Name [config.txt]: /pub/configurations/config.txt
Password: xxxxx

configUpload complete: All config parameters are uploaded.
switch:admin>
```

The following example shows **configUpload** run on a switch with Admin Domains:

```
swtich:AD5:admin> ad --select 5
Please change passwords for switch default accounts now.
Use Control-C to exit or press 'Enter' key to proceed.

Password was not changed. Will prompt again at next login
until password is changed.
swtich:AD5:admin> configUpload
Protocol (scp or ftp) [ftp]:
Server Name or IP Address [host]: 10.1.2.3
User Name [user]: JohnDoe
File Name [config.txt]: /pub/configurations/config.txt
Password: xxxxx

configUpload complete: Only zoning parameters are uploaded from ad5.
swtich:AD5:admin>
```

5. Store a soft copy of the switch configuration information in a safe place for future reference.

Note

The configuration file is printable, but you might want to see how many pages will be printed before you send it to the printer; you might not want to print a lot of pages if it is too long.

Troubleshooting Configuration Upload

If the configuration upload fails, it may be because:

- The host name is not known to the switch.
- The host IP address cannot be contacted.
- You do not have configuration upload permission on the switch.
- You do not have permission on the host to perform configuration upload.

Resolve the issue and try the command again.

Restoring Switch Information

Save the following files in soft or hard copy in a safe place for emergency reference:

Table 6-2 CLI Commands to Display Switch Configuration Information

Command	Displays
configShow	System configuration parameters and settings, including license information, zoning, and licensing information
ipAddrShow	The IP address
licenseShow	The license keys you have installed and provides better detail than the license information from the **configShow** command
fosConfig	Fabric OS features
iscsiCfg	iSCSI entities
iscsiTargetName	The IQN prefix
iscsiPortCfg	The iSCSI port parameters
isnsCfg	The configuration state of the iSNS client operation
fcLunQuery	A list of LUN IDs and LUNs for all accessible targets
portCfgEXPort	EX_Port configuration parameters
portCfgVEXPort	VEX_Port configuration parameters
fcrRouterPortCost	FC Router route information
fcrXlateConfig	A translate (xlate) domain's domain ID for both EX_Port-attached fabric and backbone fabric

Restoring a Configuration

Restoring a configuration involves overwriting the configuration on the switch by downloading a previously saved backup configuration file. Make sure that the configuration file you are downloading is compatible with your switch model, because configuration files from other model switches might cause your switch to fail.

The configuration download process is additive, that is, the lines read from the files are added to the current switch configuration. You can change a single configuration variable by downloading a file with that specific variable only. When you do so, all other variables remain unchanged.

If your setup supports anonymous users, and you log in as an anonymous user, password is still a required field, even though its value may be ignored by the FTP service.

Configuration Download Without Disabling a Switch

Starting in Fabric OS v5.2.0, you can download configuration files to a switch while the switch is enabled, that is, you do not need to disable the switch for changes in SNMP, Fabric Watch, and ACL parameters. When you use the **configDownload** command, you will be prompted to disable the switch *only when necessary.* However, if there is any changed parameter that does not belong to SNMP, Fabric Watch, or ACL, then you must disable the switch.

Configuration download without disabling a switch is independent of the hardware platform and supported on all hardware platforms running Fabric OS v5.2.0 and higher.

You may wish to first attempt downloading a configuration file first without disabling the switch. If there is at least one changed parameter outside of SNMP, Fabric Watch, and ACL, you will be prompted to disable the switch before proceeding.

To restore a configuration

1. Verify that the FTP service is running on the server where the backup configuration file is located.
2. Connect to the switch and log in as admin.
3. If there are any changed parameters in the configuration file that do no belong to SNMP, Fabric Watch, or ACL, disable the switch by entering the **switchDisable** command.
4. Enter the **configDownload** command.

   The command becomes interactive and you are prompted for the required information.

5. Respond to the prompts as follows:

Protocol (scp or ftp)	If your site requires the use of Secure Copy, specify scp. Otherwise, specify ftp.
Server Name or IP Address	Enter the name or IP address of the server where the file is stored; for example, 192.1.2.3. You can enter a server name if DNS is enabled.
User name	Enter the user name of your account on the server; for example, *JohnDoe*.
File name	Specify the full path name of the backup file; for example, */pub/configurations/config.txt*.
Password	Enter your account password for the server. Note that this is a required field even if you are logged in as an anonymous user; in such cases, the value may be ignored by the FTP service.

6. At the "Do you want to continue [y/n]" prompt, enter **y**.
7. Wait for the configuration to be restored.

   The following example shows **configDownload** run on a switch without Admin Domains:

```
switch:admin> configdownload
Protocol (scp or ftp) [ftp]: ftp
Server Name or IP Address [host]: 192.1.2.3
User Name [user]: JohnDoe
File Name [config.txt]: /pub/configurations/config.txt

               *** CAUTION ***

This command is used to download a backed-up configuration
for a specific switch. If using a file from a different
switch, this file's configuration settings will override
any current switch settings. Downloading a configuration
file, which was uploaded from a different type of switch,
may cause this switch to fail. A switch reboot might be
required for some parameter changes to take effect.

configDownload operation may take several minutes
to complete for large files.

Do you want to continue [y/n]: y
Password: xxxxx

Activating configDownload: Switch is disabled
Updated RADIUS configuration
configDownload complete: All config parameters are downloaded.
switch:admin>
```

The following example shows **configDownload** run on a switch with Admin Domains:

```
switch:AD5:admin>configdownload
Protocol (scp or ftp) [ftp]:
Server Name or IP Address [host]: 10.1.2.3
User Name [user]: JohnDoe
File Name [config.txt]: /pub/configurations/config.txt

                  *** CAUTION ***
This command is used to download a backed-up configuration
for a specific switch.  If using a file  from a  different
switch, this file's configuration  settings will  override
any current switch settings.   Downloading a configuration
file, which was uploaded  from a different type of switch,
may cause this switch to fail.  A switch reboot might be
required for some parameter changes to take effect.

configDownload operation may take several minutes
to complete for large files.

Do you want to continue [y/n]: y
Password: xxxxx

Activating configDownload: Switch is disabled

configDownload complete: Only zoning parameters are downloaded to ad5.

switch:AD5:admin>
```

8. If you disabled the switch, when the process is finished, enter the **switchEnable** command.

Note

Because some configuration parameters require a reboot to take effect, after you download a configuration file you must reboot to be sure that the parameters are enabled. Before the reboot, this type of parameter is listed in the configuration file, but it is not effective until after the reboot.

Security Considerations

Security parameters and the switch's identity cannot be changed by **configDownload**. Parameters such as the switch's name and IP address (lines in the configuration file that begin with “boot”) are ignored. Security parameters (lines in the configuration file that begin with “sec”), such as secure mode setting and version stamp, are ignored.

For more detailed information on security, see Chapter 4, “Configuring Standard Security Features”.

Troubleshooting Configuration Download

If the configuration download fails, check that these conditions are true:

- If you did not disable the switch, you must do so for some configuration downloads (see "Configuration Download Without Disabling a Switch" on page 6-5).
- The host name is known to the switch.
- The host IP address can be contacted.
- You have permission on the host to perform configuration download.
- The configuration file you are trying to download exists on the host.
- The configuration file you are trying to download is a switch configuration file.
- If you selected the (default) FTP protocol, the FTP server is running on the host;
- The configuration file uses correct syntax.

If for some reason the switch reboots during the configuration download, issue the command again.

There may be some restrictions if you are using Admin Domains. See Chapter 7, "Managing Administrative Domains" for details.

Messages Captured in the Logs

Configuration download generates both RAS log and Audit log messages resulting from execution of the **configDownload** command.

The following messages are written to the logs:

- configDownload completed successfully … (RAS and Audit log)
- configUpload completed successfully … (RAS log)
- configDownload not permitted … (Audit log)
- configUpload not permitted … (RAS log)
- (Warning) Downloading configuration without disabling the switch was unsuccessful. (Audit log)

Restoring Configurations in a FICON Environment

If the switch is operating in a FICON CUP environment, and the ASM (active=saved) bit is set on, then the switch ignores the IPL file downloaded when you restore a configuration. Table 6-3 describes this behavior in more detail.

Table 6-3 Backup and Restore in a FICON CUP Environment

ASM bit	Command	Description
on or off	**configUpload**	All the files saved in file access facility are uploaded to the management workstation. A section in the uploaded configuration file labeled FICON_CUP is in an encoded format.
on	**configDownload**	Files saved on the switch that are also present in the FICON_CUP section of the configuration file are overwritten. Files in the FICON section of configuration file that are not currently present on the switch are saved. The IPL file is not replaced, because active=saved mode is on. A warning message is displayed in the syslog to warn that the IPL file is not being overwritten.
off	**configDownload**	Files saved on the switch that are also present in the FICON_CUP section of the configuration file are overwritten. Files in the FICON section of configuration file that are not currently present on the switch are saved. The IPL file is replaced, because active=saved mode is off.

If fmsmode is enabled in a configuration file, but is disabled on the switch, the **configDownload** command fails and displays an error message. This prevents undesirable conditions that could result from enabling fmsmode on a switch that does not require it.

Downloading Configurations Across a Fabric

To save time when configuring fabric parameters and software features, you can save a configuration file from one switch and download it to other switches of the same model type, as shown in the following procedure.

Do not download a configuration file from one switch to another switch that is a different model, because it can cause the switch to fail. If you need to reset affected switches, issue the **configDefault** command.

Note

In some cases, the **configRemoveAll** root command is required for reset, in which case you should contact your system administrator or switch service provider.

To download a configuration file from one switch to another same model switch

1. Configure one switch first.
2. Use the **configUpload** command to save the configuration information. See "Backing Up a Configuration" on page 6-2.
3. First run **configDefault** on each of the target switches, and then use the **configDownload** command to download the configuration file to each of the target switches. See "Restoring a Configuration" on page 6-5.

SilkWorm 48000 Configuration Form

Use this form as a hard copy reference for your configuration information.

See the end of Chapter 7 in the *SilkWorm 48000 Hardware Reference Guide* for FC port setting tables that can be used on 16-, 32-, and 48-port blades to record configuration information.

Table 6-4 SilkWorm Configuration and Connection

SilkWorm Configuration Settings	
IP address	
Gateway address	
Chassis configuration option	
Management Connections	
Serial cable tag	
Ethernet cable tag	
Configuration Information	
Domain ID	
Switch name	
Ethernet IP address	
Ethernet subnet mask	
Total number of local devices (**nsShow**)	
Total number of devices in fabric (**nsAllShow**)	
Total number of switches in the fabric (**fabricShow**)	

Chapter

7

Managing Administrative Domains

This chapter describes the concepts and procedures for using the administrative domain feature introduced in Fabric OS v5.2.0 and contains the following topics:

- "About Administrative Domains" on page 7-1
- "Managing Admin Domains" on page 7-12
- "Using Admin Domains" on page 7-21

About Administrative Domains

An *Administrative Domain* (Admin Domain or AD) is a logical grouping of fabric elements that defines what switches, ports, and devices you can view and modify. An Admin Domain is a filtered administrative view of the fabric.

Note

If you do not implement Admin Domains, the feature has no impact on users and you do not need to learn how to use this functionality.

Admin Domains permit access to a configured set of users. Using Admin Domains, you can partition the fabric into logical groups and allocate administration of these groups to different user accounts so that these accounts manage only the Admin Domains assigned to them and do not make changes to the rest of the fabric.

For example, you can put all the devices in a particular department in the same Admin Domain for ease of managing those devices. If you have remote sites, you could put the resources in the remote site in an Admin Domain and assign the remote site administrator to manage those resources.

You set up zones to define which devices and hosts can communicate with each other; you set up Admin Domains to define which users can manage which devices, hosts, and switches.

You can have up to 256 Admin Domains in a fabric (254 user-defined and 2 system-defined), numbered from 0 through 255. Admin Domains are designated by a name and a number. This document refers to specific Admin Domains using the format "AD*n*" where *n* is a number between 0 and 255.

Note

Do not confuse an *Admin Domain number* with the *domain ID of a switch*. They are two different identifiers.

Before using the procedures described in this chapter, you should become familiar with the Admin Domain concepts described in the following sections.

An “AD-aware switch” is a switch that runs Fabric OS v5.2.0 or higher (on both CPs, if a dual CP switch) and has a valid Advanced Zoning license.

An “AD-unaware switch” is a switch that is running:

- Fabric OS v5.1.x or lower
- Fabric OS v5.2.0 or higher but does not have an Advanced Zoning license
- Fabric OS v5.2.0 or higher on one CP but Fabric OS v5.1.x or lower on the other (for dual-CP switches) and is in HA_Sync

Figure 7-1 shows a fabric with two Admin Domains: AD1 and AD2.

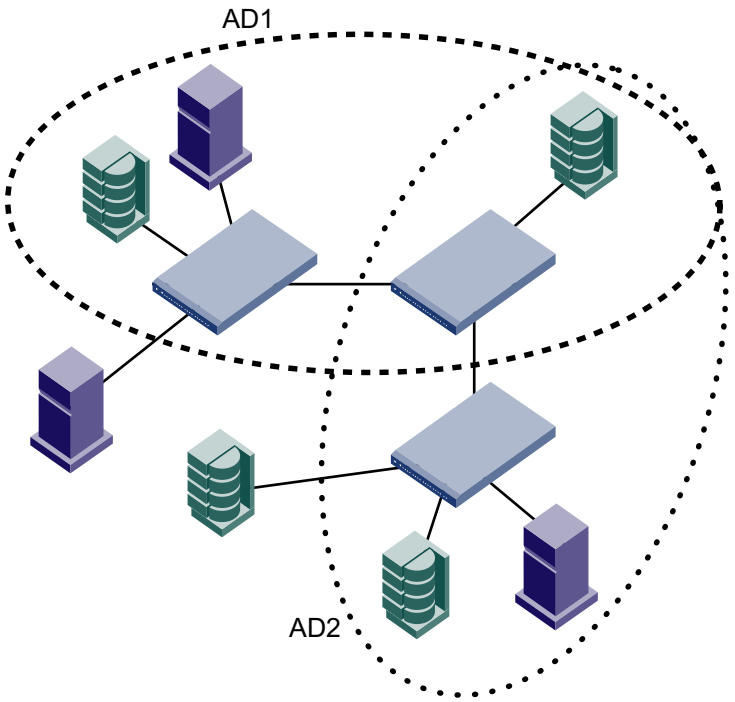

Figure 7-1 Fabric with Two Admin Domains

Figure 7-2 shows how users get a filtered view of this fabric, depending on which Admin Domain they are in. As shown in Figure 7-2, users can see all switches and E_Ports in the fabric, regardless of their Admin Domain; however, the switch ports and end devices are filtered based on Admin Domain membership.

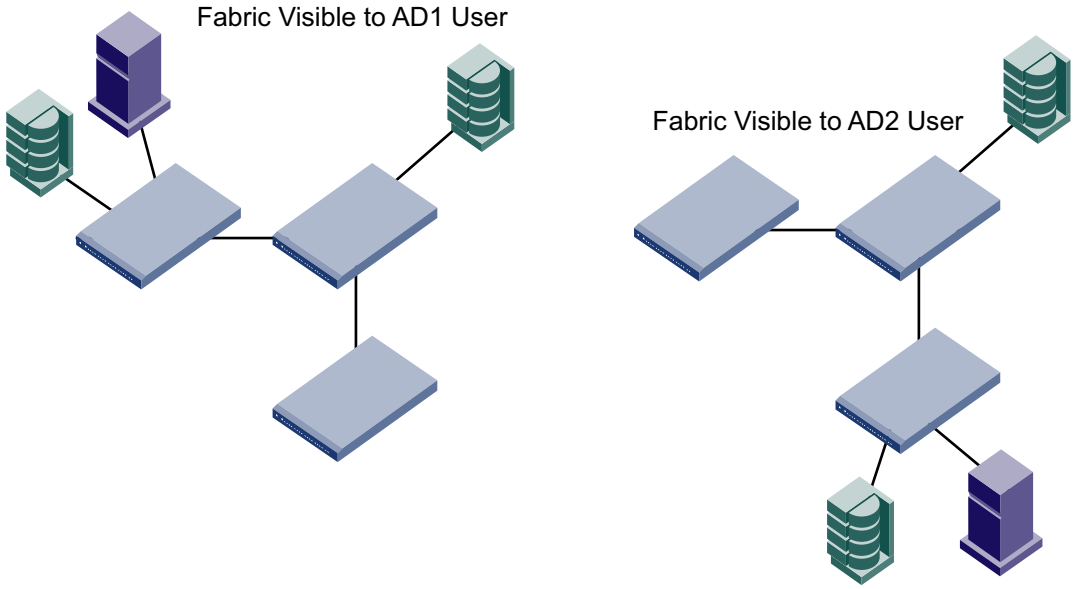

Figure 7-2 Filtered Fabric Views

Admin Domain Features

Admin Domains allow you to:

- Define the scope of an Admin Domain to encompass ports and devices within a switch or a fabric.
- Share resources across multiple Admin Domains. For example, you can share array ports and tape drives between multiple departments. In Figure 7-1, one of the storage devices is shared between AD1 and AD2.
- Have a separate zone database for each Admin Domain. See "Admin Domains, Zones, and Zone Databases" on page 7-26 for more information.
- Move devices from one Admin Domain to another without traffic disruption, cable reconnects, or discontinuity in zone enforcement.
- Provide strong fault and event isolation between Admin Domains.
- Have visibility of all physical fabric resources. All switches, E_Ports, and FRUs (including blade information) are visible.
- Implement Admin Domains in a fabric with some switches running AD-unaware firmware versions (that is, firmware versions lower than Fabric OS v5.2.0).
- Continue to run existing third-party management applications. Prior and existing versions of third party management applications continue to work with admin and user IDs.

Note

Fabric OS v5.2.0 supports up to 16 Admin Domains in a fabric. Admin Domain capacity has been tested for up to 254 Admin Domains per fabric, although greater than 16 Admin Domains is not supported.

Requirements for Admin Domains

This section lists the requirements for implementing Admin Domains in a fabric. For information about supported AD platforms, see the *Fabric OS v5.2.0 Release Notes*.

- Admin Domains are supported only on fabrics with one or more switches running Fabric OS 5.2.0 and higher.
- You must have a valid Advanced Zoning license to use Admin Domains.
- The default zone mode setting must be set to No Access before you create Admin Domains (see "Implementing Admin Domains" on page 7-13).
- To use Admin Domains and the FC-FC Routing Service in the same fabric, the switches connecting to the FC Router must be running Fabric OS 5.2 or higher.
- If you are using LSAN zones:
  - Do not use LSAN zone names ending with "_AD*n*" (where *n* is the Admin Domain number).
  - Do not use LSAN zone names longer than 57 characters.
- You must be in the native operating mode to use Admin Domains. Admin Domains are not supported in interoperability mode.
- Gigabit Ethernet (GbE) ports cannot be members of an Admin Domain.

User-Defined Administrative Domains

AD1 through AD254 are user-defined Admin Domains. These user-defined Admin Domains can be created only by a physical fabric administrator (see "Admin Domain Access Levels" on page 7-6 for more information).

In Figure 7-1, AD1 and AD2 are user-defined Admin Domains.

System-Defined Administrative Domains

When you install Fabric OS v5.2.0 firmware, the switch enters AD-capable mode with domains AD0 and AD255 automatically created. AD0 and AD255 are special Admin Domains. AD0 and AD255 always exist and cannot be deleted or renamed. They are reserved for use in creation and management of Admin Domains.

AD0

AD0 is a system-defined Admin Domain that, in addition to containing members you explicitly added (similar to user-defined Admin Domains), contains all online devices, switch ports, and switches that have not been assigned to any user-defined Admin Domain.

Unlike user-defined Admin Domains, AD0 has an implicit and an explicit membership list. User-defined Admin Domains have only explicit members.

- The "implicit membership list" contains all devices, switch ports, and switches that have not been assigned to any other Admin Domain.

Initially, the AD0 implicit membership list contains all devices, switch ports, and switches in the fabric. This implicit membership list includes all AD-unaware switches and the devices attached to them. When you explicitly create AD1 through AD254, the devices, switch ports, and switches used to create these user-defined Admin Domains disappear from the AD0 implicit membership list.

- The "explicit membership list" contains all devices, switch ports, and switches that you explicitly add to AD0 and can be used to force device and switch sharing between AD0 and other Admin Domains. Explicit members can be added to AD0 using the **ad --add** command.

AD0 can be managed like any user-defined Admin Domain. The only difference between AD0 and user-defined Admin Domains is the implicit membership list.

The implicit members of AD0 change dynamically as the membership of other Admin Domains changes. The explicit members of AD0 are not deleted unless you explicitly remove them.

For example, if you explicitly add DeviceA to AD0 and it is not a member of any other Admin Domain, then DeviceA is both an implicit and an explicit member of AD0. If you add DeviceA to AD2, then DeviceA is deleted from the AD0 implicit membership list, but is *not* deleted from the AD0 explicit membership list. If you then remove DeviceA from AD2, DeviceA is added back to the AD0 implicit membership list (assuming DeviceA is not in any other Admin Domain).

When a new device is added to the fabric, it automatically becomes an implicit member of AD0 until it is explicitly added to an Admin Domain.

AD0 is useful when you create Admin Domains because you can see which devices, switch ports, and switches have not yet been assigned to any Admin Domains.

AD0 owns the root zone database (legacy zone database). During zone merge or zone update, only the root zone database is exchanged with AD-unaware switches.

AD255

AD255 is used for Admin Domain management. You can use AD255 to get an unfiltered view of the fabric and to view the hierarchical zone databases of AD0 through AD254. All Admin Domain management is done in the AD255 context.

AD255 does not have a zone database associated with it; you cannot use AD255 to perform any zoning management tasks (non-read operations such as creating or modifying zones).

Figure 7-3 shows the same fabric from Figure 7-1, but with AD0 and AD255 shown. AD0 contains the two devices that are not in any of the user-defined Admin Domains (AD1 and AD2). AD255 encompasses the entire physical fabric.

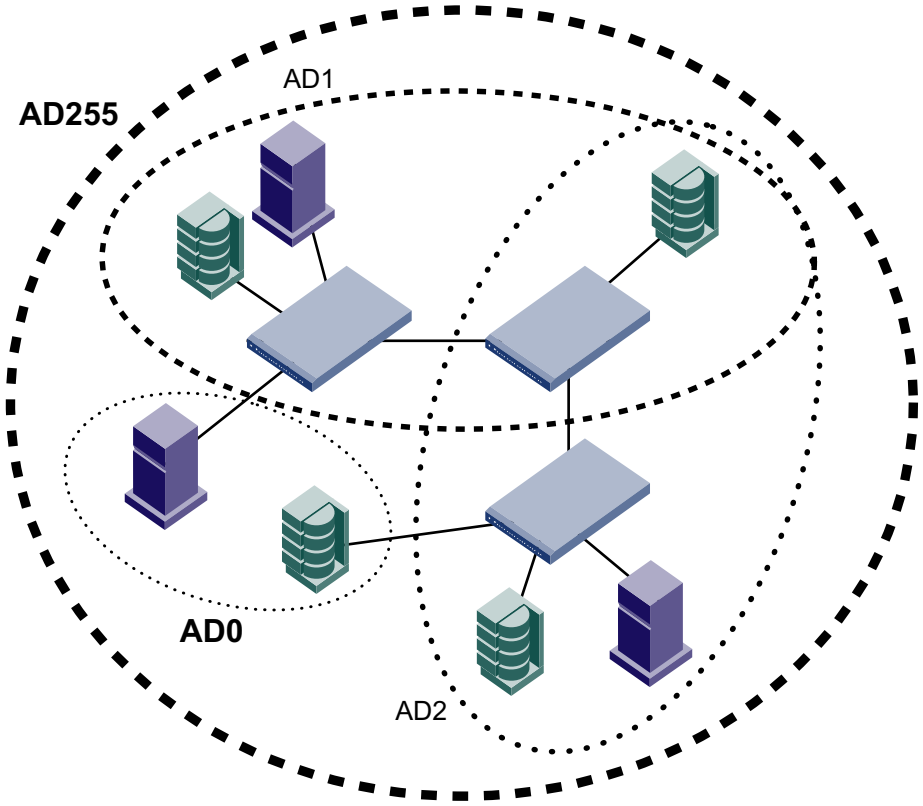

Figure 7-3 Fabric with AD0 and AD255

Admin Domain Access Levels

Admin Domains offer a hierarchy of administrative access.

To manage Admin Domains, you must be a physical fabric administrator. A “physical fabric administrator” is a user with the Admin role and access to all Admin Domains (AD0 through AD255).

Other administrative access is determined by your defined RBAC role and AD membership. Your role determines your access level and permission to perform an operation. Your AD membership determines the fabric resources that you can operate on.

Table 7-4 lists each Admin Domain user type and describes its administrative access and capabilities.

Table 7-4 AD User Types

User Type	Description
Physical Fabric Administrators	• User account with Admin role and with access to all Admin Domains (AD0 through AD255). • Create and manage all Admin Domains. Only a physical fabric administrator can perform Admin Domain configuration and management. • Assign other administrators or users to each Admin Domain. • Only a physical fabric administrator can create other physical fabric administrators.
Administrative Domain Users	• Can be assigned to one or more Admin Domains. • Manage the resources within their Admin Domains. • If their role permits, can create user accounts and assign them to Admin Domains in their list. • Cannot view other Admin Domain definitions. They can view only members of their own Admin Domains.

Admin Domains and Login

You are always logged in to an Admin Domain, and you can view and modify only the devices in that Admin Domain.

If you have access to more than one Admin Domain, one of them will have been specified as your "home Admin Domain," the one you are automatically logged in to. If your home Admin Domain is deleted or deactivated, then by default you are logged in to the lowest numbered active Admin Domain in your Admin Domain List. The home Admin Domain, like the Admin Domain list, is a configurable property of a non-default user account. Here is some information about AD accounts:

- You can log in to only one Admin Domain at a time. You can later change the Admin Domain to which you are logged in (see "Switching to a Different Admin Domain Context" on page 7-23).
- For default accounts such as admin and user, the home Admin Domain defaults to AD0 and cannot be changed.
- The Admin Domain list for the default admin account is 0–255, which gives this account automatic access to any Admin Domain as soon as the domain is created, and makes this account a physical fabric administrator.
- The Admin Domain list for the default user account is AD0 only.
- For user-defined accounts, the home Admin Domain also defaults to AD0 but an administrator can set the home Admin Domain to any Admin Domain to which the account has been given access.
- If you are in any Admin Domain context other than AD0, the Admin Domain is included in your prompt. The following are example prompts for when you are in the AD0, AD1, and AD255 contexts, respectively:

```
sw5:admin>
sw5:AD1:admin>
sw5:AD255:admin>
```

Admin Domain Member Types

You define an Admin Domain by identifying members of that domain. Admin Domain members can be devices, switch ports, or switches. Defining these member types is similar to defining a traditional zone member type. An Admin Domain does not require or have a new domain ID or management IP address linked to it.

The following sections describe these member types in more detail.

Device Members

Device members are defined by the device World Wide Name (WWN). A device member:

- Can be either a device port WWN or device node WWN.
- Can be a physical or virtual device, such as an FCR Proxy device.
- Grants view access to the device and zoning rights.
- Provides a pure virtual view. The cabling and switch port diagnostics and control is done by the physical fabric administrator.

View rights are also granted to the switch port to which the device is attached.

Port control is provided only through switch port membership and is not provided for device members. When you create an Admin Domain, the end device members do not have to be online, even though their WWNs are used in the Admin Domain definition. The physical fabric administrator performs cabling and switch port diagnostics and control.

You can share device members across multiple Admin Domains. You can also zone shared devices differently in each Admin Domain. A device WWN member does not automatically grant usage of corresponding (domain, port) members in the zone configuration. If you specify a device WWN member in the Admin Domain member list, zone enforcement ignores zones with the corresponding port (the port to which the device is connected) member usage.

Switch Port Members

Switch port members are defined by switch (domain, port). A switch port member:

- Grants port control rights and zoning rights for that switch port.
- Grants view access and zoning rights to the device connected to that switch port.
- Allows you to share (domain, port) members across multiple Admin Domains. In each Admin Domain, you can also zone shared devices differently.
- Implicitly includes all devices connected to the specified (domain, port) members in the Admin Domain membership.
- Allows you to specify a range of ports as Admin Domain members. For example: `<D,[0-15]>`. The port range arguments are expanded and stored in the Admin Domain member list.

If a device is a member of an Admin Domain, the switch port to which the device is connected becomes an indirect member of that Admin Domain and the (domain, port) is removed from the AD0 implicit membership list.

Note

The (domain, port) members are not automatically changed when the switch domain ID changes.

Switch Members

Switch members are defined by the switch WWN or domain ID. A switch member:

- Grants administrative control to the switch.
- Grants port control for all ports in that switch.
- Allows switch administrative operations such as **switchDisable**, **switchEnable**, **reboot**, and **firmwareDownload**.
- Does not provide zoning rights for the switch ports or devices.

To allow devices to be zoned within Admin Domains, you must specify the port members using (domain, port) or device WWN members.

E_Ports (E_Ports, VE_Ports, EX_Ports, and VEX_Ports) are implicitly shared across all Admin Domains. An administrator can perform port control (such as **portDisable** or **portEnable**) only if the (domain, port) of the E_Port is part of the Admin Domain.

Admin Domains and Switch WWN

Admin Domains are treated as fabrics. Because switches cannot belong to more than one fabric, switch WWNs are converted so that they appear as unique entities in different Admin Domains (fabrics). This WWN conversion is done only in the AD1 through AD254 context. AD0 and AD255 use unconverted switch WWNs.

The switch WWN is in the following format:

10:00:*nn*:*nn*:*nn*:*nn*:*nn*:*nn*

In an Admin Domain context, the switch WWN is converted from NAA=1 to NAA=5 format, with the Admin Domain number added, using the following syntax:

5*n*:*nn*:*nn*:*nn*:*nn*:*nn*:*n*9:*xx*

where *xx* is the AdminDomain_number.

For example, if the switch WWN is:

10:00:00:60:69:e4:24:e0

then the converted WWN for that switch in AD1 would be:

50:06:06:9e:42:4e:09:01

Figure 7-5 shows an unfiltered view of a fabric with two switches, three devices, and two Admin Domains. The devices are labeled with device WWN and the switches are labeled with domain ID and switch WWN.

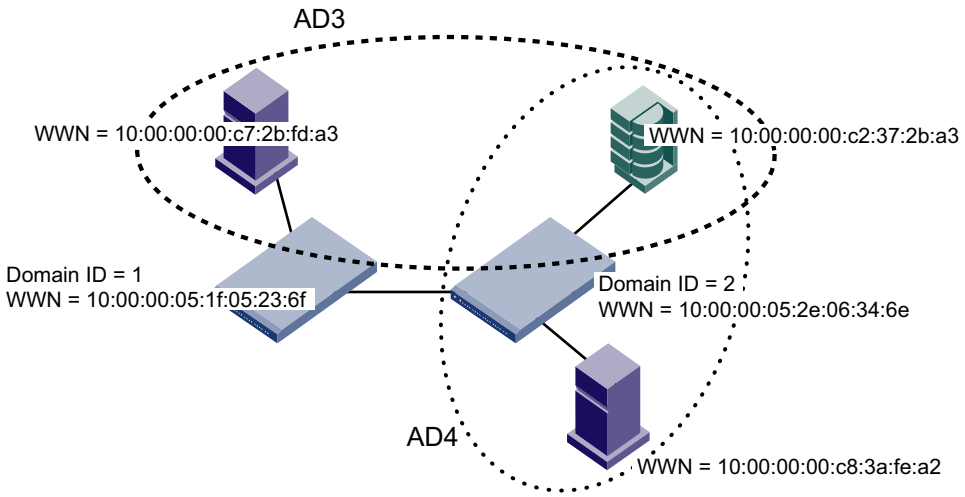

Figure 7-5 Fabric Showing Switch and Device WWNs

Figure 7-6 shows the filtered view of the fabric as seen from AD3 and AD4. The switch WWNs are converted to the NAA=5 syntax; the device WWNs and domain IDs remain the same.

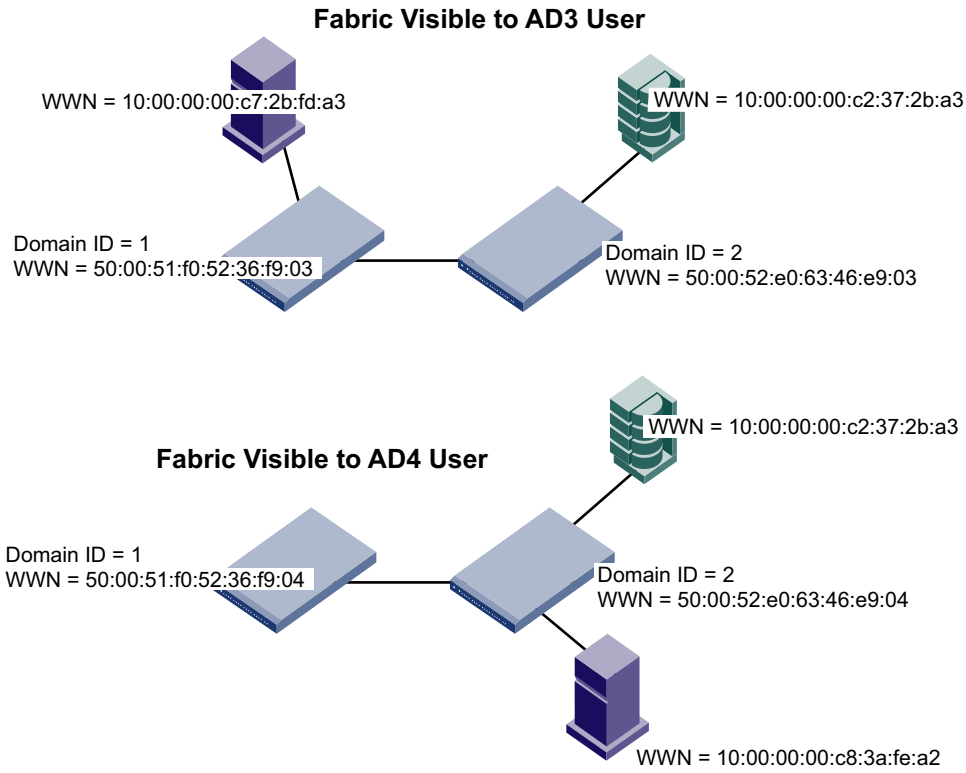

Figure 7-6 Filtered Fabric Views Showing Converted Switch WWNs

Admin Domain Compatibility and Availability

Admin Domains maintain continuity of service for Fabric OS v5.2.0 features and operate in mixed-release fabric environments. High availability is supported along with some backward compatibility. The following sections describe the continuity features of Admin Domain usage.

Admin Domains and Merging

When an E_Port comes online, the adjacent switches merge their AD databases. The receiving switch accepts an AD database from the neighboring switch only if the local AD database is empty or if the new AD database exactly matches both the defined and effective configurations of the local AD database. If the AD database merge fails, the E_Port is segmented with “AD conflict” error code.

Compatibility

Admin Domains can be implemented in fabrics with mix of AD-aware switches and AD-unaware switches. The following considerations apply:

- In mixed-fabric configurations, the legacy switches allow unfiltered access to the fabric and its devices; hence, these legacy switches should be managed by the physical fabric administrator.
- You must zone all ports and devices from legacy switches in the AD0 root zone database.
- If you have legacy switches in your AD-activated fabric, you must ensure that all new AD resources have enough interconnectivity so that they do not get isolated into subfabrics with a legacy subfabric interposed in the middle, as shown in Figure 7-7.

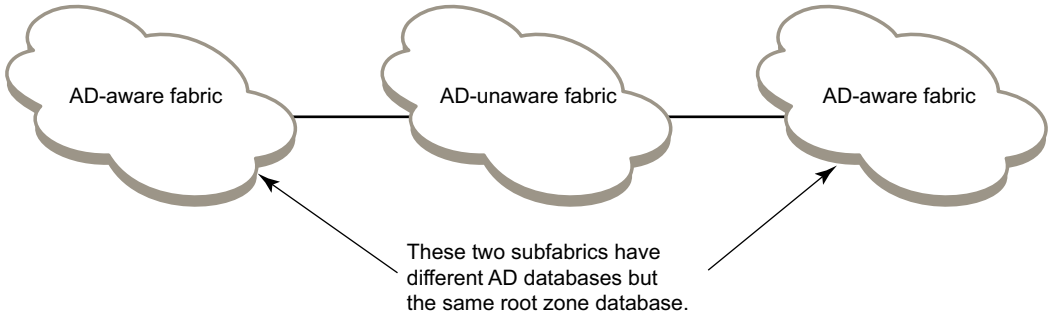

Figure 7-7 Isolated Subfabrics

Firmware Upgrade and Downgrade Scenarios

You cannot perform a firmware downgrade from Fabric OS v5.2.0, if ADs are configured in the fabric. Following are special scenarios for director class products only:

- If the primary and secondary CPs are running pre-Fabric OS v5.2.0 and are in HA-Sync, if **firmwaredownload** is used to upgrade one CP alone (using the **firmwaredownload -s** option), then that CP will run in an AD-unaware mode (AD creation operations will fail and the local switch will show up as an AD-unaware switch in the fabric).
- If the primary and secondary CPs are running Fabric OS v5.2.0 and if ADs are configured, any attempt to downgrade one or both CPs to pre Fabric OS v5.2.0 versions will fail. The Admin Domain configuration must be cleared before you can perform the downgrade (see “Deleting All User-Defined Admin Domains” on page 7-19).

Managing Admin Domains

This section is for physical fabric administrators who are managing Admin Domains. You must be a physical fabric administrator to perform the tasks in this section.

- "Implementing Admin Domains" on page 7-13
- "Creating an Admin Domain" on page 7-13
- "Assigning a User to an Admin Domain" on page 7-15
- "Activating and Deactivating Admin Domains" on page 7-16
- "Adding and Removing Admin Domain Members" on page 7-17
- "Renaming an Admin Domain" on page 7-18
- "Deleting an Admin Domain" on page 7-19
- "Deleting All User-Defined Admin Domains" on page 7-19
- "Validating an Admin Domain Member List" on page 7-20

Understanding the AD Transaction Model

You use the **ad** command to perform most of the tasks in this section. This command follows a batched-transaction model, which means that changes to the Admin Domain configuration occur in the transaction buffer.

An Admin Domain configuration can exist in several places:

- **Effective configuration—**The Admin Domain configuration that is currently in effect.
- **Defined configuration—**The Admin Domain configuration that is saved in flash memory. There might be differences between the effective configuration and the defined configuration.
- **Transaction buffer—**The Admin Domain configuration that is in the current transaction buffer and has not yet been saved or canceled.

How you end the transaction determines the disposition of the Admin Domain configuration in the transaction buffer. The following commands end the Admin Domain transaction:

ad --save	Saves the changes in the transaction buffer to the defined configuration in persistent storage and propagates the defined configuration to all switches in the fabric. Note that for delete and clear operations, if one or more of the deleted Admin Domains are in the effective configuration, you cannot use **--save**, but must use **--apply** instead.
ad --apply	Saves the changes to the defined configuration in persistent storage and enforces the defined configuration on all switches in the fabric, replacing the effective configuration.
ad --transabort	Aborts the transaction and clears the transaction buffer. The effective and defined configurations remain unchanged.

You can enter the **ad --transshow** command at any time to display the ID of the current Admin Domain transaction.

Detailed information about CLI syntax and options is available in the *Fabric OS Command Reference.*

Implementing Admin Domains

To begin implementing an Admin Domain structure within your SAN, you must first set the default zone mode to No Access. You must be in AD0 to change the default zone mode. You can use the **defZone --show** command to see the current default zone mode setting.

To set the default zone mode

1. Log in to an AD-aware switch in the fabric with the appropriate RBAC role.
2. Ensure you are in the AD0 context.

   Use the **ad --show** command to determine the current Admin Domain and the **ad --select 0** command to switch to the AD0 context, if necessary.
3. Enter the **defZone --noAccess** command, and respond to the prompt with a **y**.
4. Enter the **cfgSave** command.

```
sw5:admin> defzone --noaccess
You are about to set the Default Zone access mode to No Access
Do you want to set the Default Zone access mode to No Access ? (yes, y, no, n):
[no] y

sw5:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only?  (yes, y, no, n): [no] y
Updating flash ...
sw5:admin>
```

Creating an Admin Domain

To create an Admin Domain, you must specify an Admin Domain name or number or both.

- If you specify only a number, the Admin Domain name is automatically assigned to be "AD*n*", where *n* is the number you specified.
- If you specify only a name, the Admin Domain number is automatically assigned to be the lowest available AD number, except if you specify a name in the format "AD*n*", in which case the Admin Domain number is assigned to be *n*.

When you create an Admin Domain, you must specify at least one member (switch, switch port, or device). You cannot create an empty Admin Domain. For more information about these member types, see "Admin Domain Member Types" on page 7-8.

You create Admin Domains in the transaction buffer using the **ad --create** command. You can either save the newly created Admin Domain to a defined configuration (using **ad --save**) or make it the effective Admin Domain configuration directly (using **ad --apply**).

The following procedures describe the steps for creating Admin Domains and include examples.

Before creating an Admin Domain, you must set the default zone mode to No Access using the **defZone** command, as described in "Implementing Admin Domains" on page 7-13.

To create an Admin Domain

1. Log in as the physical fabric administrator to an AD-aware switch in the fabric.
2. Switch to the AD255 context, if you are not already in that context:

   ```
   ad --select 255
   ```

3. Enter the **ad --create** command using the **-d** option to specify device and switch port members and the **-s** option to specify switch members:

   ```
   ad --create ad_id -d "dev_list" -s "switch_list"
   ```

   where *ad_id* is the Admin Domain name or number, *dev_list* is a list of device WWNs or (domain, port) members, and *switch_list* is a list of switch WWNs or domain IDs.

4. *Optional:* To end the transaction now, enter **ad --save** to save the Admin Domain definition or enter **ad --apply** to save the Admin Domain definition and directly apply the definitions to the fabric.

   The following example creates Admin Domain AD1, consisting of two switches, which are designated by domain ID and switch WWN.

   ```
   sw5:AD255:admin> ad --create AD1 -s "97; 10:00:00:60:69:80:59:13"
   ```

   The following example creates Admin Domain “blue_ad,” consisting of two switch ports (designated by (domain, port)), one device (designated by device WWN), and two switches (designated by domain ID and switch WWN).

   ```
   sw5:AD255:admin> ad --create blue_ad –d "100,5; 1,3; 21:00:00:e0:8b:05:4d:05; –s
   "97; 10:00:00:60:69:80:59:13"
   ```

Assigning a User to an Admin Domain

After you create an Admin Domain, you can specify one or more user accounts as the valid accounts who can use that Admin Domain. You create these user accounts using the **userConfig** command. User accounts have the following characteristics with regard to Admin Domains:

- A user account can only have a single role.

  You can choose roles from the one of the seven types of roles, either the existing user and administrator role or one of the other RBAC roles.

- You can configure a user account to have access to the physical fabric through AD255 and to a list of Admin Domains (AD0–AD254).
- You can configure a user account to have access to only a subset of your own Admin Domain list. Only a physical fabric administrator can create another physical fabric administrator user account.
- Users capable of using multiple Admin Domains, can designate one of these Admin Domains as the home Admin Domain, which is the default Admin Domain context after login.
- If you do not specify one, the home Admin Domain is the lowest valid Admin Domain in the numerically-sorted AD list.
- Users can log in to their Admin Domains and create their own Admin Domain-specific zones and zone configurations.
- Adding an Admin Domain list, home Admin Domain, and role to a user configuration is backward compatible with pre-Fabric OS v5.2.0 firmware. When you downgrade to pre-Fabric OS v5.2.0 firmware, the **userConfig** command records are interpreted using legacy logic.

To create a new user account for managing Admin Domains

1. Connect to the switch and log in as admin.
2. Enter the **userconfig --add** command using the **-r** option to set the role, the **-a** option to provide access to Admin Domains, and the **-h** option to specify the home Admin Domain.

   userconfig --add *username* **-r** *role* **-h** *home_AD* **-a "***AD_list***"**

   where *username* is the name of the account, *role* is the user account role, *home_AD* is the home Admin Domain, and *AD_list* is the list of Admin Domains to which the user account will have access.

   The following example creates new user account ad1admin with an admin role and assigns one Admin Domain, blue_ad1, to it. This example also assigns blue_ad1 as the user's home Admin Domain.

   ```
   sw5:admin> userconfig --add ad1admin -r admin -h blue_ad1 -a "blue_ad1"
   ```

   The following example creates new user account ad2admin with an admin role, access to Admin Domains 1 and 2, and home Admin Domain set to 2.

   ```
   sw5:admin> userconfig --add ad2admin -r admin -h 2 -a "1,2"
   ```

To assign Admin Domains to an existing user account

1. Connect to the switch and log in as admin.
2. Enter the **userConfig --addad** command using the **-a** option to provide access to Admin Domains and the **-h** option to specify the home Admin Domain.

   userconfig --addad *username* **-h** *home_AD* **-a "***AD_list***"**

   where *username* is the name of the account, *home_AD* is the home Admin Domain, and *AD_list* is the list of Admin Domains to which the user account will have access.

   The following example assigns Admin Domain green_ad2 to the existing user account ad1admin.

```
sw5:admin> userconfig --addad ad1admin -r admin -a "green_ad2"
```

To create a new physical fabric administrator user account

1. Connect to the switch and log in as admin.
2. Enter the **userconfig --add** command using the **-r** option to set the role to admin and the **-a** option to provide access to Admin Domains 0 through 255.

   userconfig --add *username* **-r** admin **-h** *home_AD* **-a "**0-255**"**

   where *username* is the name of the account and *home_AD* is the home Admin Domain.

   The following example creates new user account pf_admin1 with an admin role, access to all Admin Domains (AD0 through AD255), and home Admin Domain set to 255. This user account is now a physical fabric administrator.

```
sw5:admin> userconfig --add pf_admin1 -r admin -h 255 -a "0-255"
```

Activating and Deactivating Admin Domains

An Admin Domain can be in either an active or inactive state. When you create an Admin Domain, it is automatically in the active state.

If you deactivate an Admin Domain, the members assigned to the Admin Domain can no longer access their hosts or storage unless those members are part of another Admin Domain.

You cannot log in to an Admin Domain that has been deactivated. You must activate an Admin Domain before you can log in to it.

To activate an Admin Domain

1. Connect to the switch and log in as admin.
2. Switch to the AD255 context, if you are not already in that context.

   ad --select 255

3. Enter the **ad --activate** option. The activate option prompts for confirmation.

   On default, after the Admin Domain is activated, the devices specified under that AD are not able to see each other until they are zoned together.

4. *Optional:* To end the transaction now, enter **ad --save** to save the Admin Domain definition or enter **ad --apply** to save the Admin Domain definition and directly apply the definitions to the fabric.

   The following example activates Admin Domain AD_B5.

```
sw5:AD255:admin> ad --activate AD_B5
You are about to activate a new admin domain.
Do you want to activate 'AD_B5' admin domain (yes, y, no, n): [no]: y
sw5:AD255:admin>
```

To deactivate an Admin Domain

1. Connect to the switch and log in as admin.
2. Disable the zone configuration under the Admin Domain you want to deactivate.
3. Switch to the AD255 context, if you are not already in that context.

   ```
   ad --select 255
   ```

4. Enter the **ad --deactivate** option.

   The **ad --deactivate** option prompts for confirmation.

5. *Optional:* To end the transaction now, enter **ad --save** to save the Admin Domain definition or enter **ad --apply** to save the Admin Domain definition and directly apply the definitions to the fabric.

   All active user sessions associated with the Admin Domain are terminated. The deactivate option does not disable ports.

   The following example deactivates Admin Domain AD_B4.

```
sw5:AD255:admin> ad --deactivate AD_B4
You are about to deactivate an AD.
This operation will fail if an effective zone configuration exists in the AD
Do you want to deactivate 'AD_B5' admin domain (yes, y, no, n): [no] y
sw5:AD255:admin>
```

Adding and Removing Admin Domain Members

Use the following procedures to add or remove members of an Admin Domain.

Note

If you remove the last member of an Admin Domain, that Admin Domain is automatically deleted.

To add members to an existing Admin Domain

1. Connect to an AD-aware switch and log in as admin.
2. Switch to the AD255 context, if you are not already in that context.

   ```
   ad --select 255
   ```

3. Enter the **ad --add** command using the **-d** option to specify device and switch port members and the **-s** option to specify switch members:

   ad --add *ad_id* -d "*dev_list*" -s "*switch_list*"

   where *ad_id* is the Admin Domain name or number, *dev_list* is a list of device WWNs or (domain, port) members, and *switch_list* is a list of switch WWNs or domain IDs.

4. *Optional:* To end the transaction now, enter **ad --save** to save the Admin Domain definition or enter **ad --apply** to save the Admin Domain definition and directly apply the definitions to the fabric.

   The following example adds two switch ports, designated by (domain, port), to Admin Domain AD1.

   ```
   sw5:AD255:admin> ad --add AD1 -d "100,5; 4,1"
   ```

To remove members from an Admin Domain

1. Connect to the switch and log in as admin.
2. Switch to the AD255 context, if you are not already in that context.

   ad --select 255

3. Enter the **ad --remove** command using the **-d** option to specify device and switch port members and the **-s** option to specify switch members:

   ad --remove *ad_id* -d "*dev_list*" -s "*switch_list*"

   where *ad_id* is the Admin Domain name or number, *dev_list* is a list of device WWNs or (domain, port) members, and *switch_list* is a list of switch WWNs or domain IDs.

   Removing the last member element of an Admin Domain deletes the Admin Domain.

4. *Optional:* To end the transaction now, enter **ad --save** to save the Admin Domain definition or enter **ad --apply** to save the Admin Domain definition and directly apply the definitions to the fabric.

   The following example removes port 5 of domain 100 and port 3 of domain 1 from AD1.

   ```
   sw5:AD255:admin> ad --remove 1 –d "100,5; 1,3"
   ```

   The following example removes switch 100 from the membership list of AD4.

   ```
   sw5:AD255:admin> ad --remove 4 –s "100"
   ```

Renaming an Admin Domain

Use this procedure if you want to change the name of an Admin Domain. You can also change auto-assigned names (AD*n*).

To rename an Admin Domain

1. Connect to the switch and log in as admin.
2. Switch to the AD255 context, if you are not already in that context.

   ad --select 255

3. Enter the **ad --rename** command with the present name and the new name.

   ad --rename *present_name new_name*

   The rename operation does not take effect if the Admin Domain you want to rename is part of the effective configuration and thus enforced.

4. *Optional:* To end the transaction now, enter **ad --save** to save the Admin Domain definition or enter **ad --apply** to save the Admin Domain definition and directly apply the definitions to the fabric.

   The Admin Domain numbers remain unchanged after the operation.

   The following example changes the name of Admin Domain Eng_AD to Eng_AD2.

```
sw5:AD255:admin> ad --rename Eng_AD Eng_AD2
```

Deleting an Admin Domain

When you delete an Admin Domain its devices no longer have access to the members of the zones with which it was associated.

To delete an Admin Domain

1. Connect to the switch and log in as admin.
2. Delete the zone database under the Admin Domain.
3. Switch to the AD255 context.

   ad --select 255

4. Enter the **ad --delete** command.

   ad --delete *ad_id*

   The **ad --delete** command prompts you for confirmation before triggering the deletion. The command will succeed, whether the administration domain is in an activated or deactivated state.

5. Enter the **ad --apply** command to save the Admin Domain definition and directly apply the definitions to the fabric.

   The following example deletes Admin Domain AD_B3.

```
sw5:AD255:admin> ad --delete AD_B3
You are about to delete an AD.
This operation will fail if zone configuration exists in the AD
Do you want to delete 'AD_B3' admin domain (yes, y, no, n): [no] y
sw5:AD255:admin>
```

Deleting All User-Defined Admin Domains

When you clear the Admin Domain configuration, all user-defined Admin Domains are deleted, the explicit membership list of AD0 is cleared, and all fabric resources (switches, ports, and devices) are returned to the implicit membership list of AD0.

You cannot clear the Admin Domain configuration if zone configurations exist in any of the user-defined Admin Domains.

To clear all Admin Domain definitions

1. Clear all individual AD zone databases, in separate transactions, before proceeding with this operation.
2. Connect to the switch and log in as admin.
3. Switch to the AD255 context, if you are not already in that context.

   ```
   ad --select 255
   ```

4. Enter the **ad --clear** command.

   This option prompts you for confirmation before triggering the delete of all Admin Domains.

5. Enter the **ad --apply** command to save the Admin Domain definition and directly apply the definitions to the fabric.

```
sw5:AD255:admin> ad --clear
You are about to delete all ADs definitions.
This operations will fail if zone configurations exists in AD1-AD254
Do you want to clear all admin domains (yes, y, no, n): [no] y
sw5:AD255:admin>
```

Validating an Admin Domain Member List

The **ad --validate** option allows you to validate the device and switch member list and flag all resources that are from AD-unaware switches. You can use the **validate** option to list Admin Domain members from AD-unaware switches and non-existing or offline Admin Domain members.

You can use the **validate** option to identify misconfigurations of the Admin Domain. For example, in fabrics with a mix of AD-aware and AD-unaware switches, elements in the Admin Domain member list from old AD-unaware switches are not enforced.

The Admin Domain validation process is not applicable for AD0, as AD0 implicitly assumes all unassigned and AD-unaware online switches and their devices.

To list the switches and devices in an AD member list

1. Connect to the switch and log in as admin.
2. Switch to the AD255 context, if you are not already in that context.

   ```
   ad --select 255
   ```

3. Enter the **ad --validate** command.

   ```
   ad --validate ad_id -m mode
   ```

   If you do not specify any parameters, the entire AD database (transaction buffer, defined configuration, and effective configuration) is displayed.

   If you do not specify an Admin Domain, information about all existing Admin Domains is displayed.

The **-m** mode flag can be used as follows:

- 0 to display the Admin Domain configuration in the current transaction buffer.
- 1 to display the Admin Domain configuration stored in the persistent memory (defined configuration).
- 2 to display the currently enforced Admin Domain configuration (effective configuration).

The following example validates the member list of Admin Domain 10 in the current transaction buffer.

```
sw5:AD255:admin> ad --validate 10 -m 0
```

Using Admin Domains

This section is for users and administrators and describes how you use Admin Domains. If you are a physical fabric administrator and you want to create, modify, or otherwise manage Admin Domains, see "Managing Admin Domains" on page 7-12.

The Admin Domain looks like a virtual switch or fabric to a user. However, based on the user role and type (User_ID), users are presented with only their relevant AD-based views (see Figure 7-1 and Figure 7-2). Any devices and switch ports that are not defined as part of the Admin Domain are not shown and are not available to that AD user.

Each Admin Domain can also have its own zone configurations (defined and effective) with zones and aliases under them.

Using CLI Commands in an AD Context

The CLI command input arguments are validated against the AD member list; they do not work with input arguments that specify resources that are not members of the current Admin Domain. All commands present filtered output—showing only the members of the current Admin Domain.

For example, **switchShow** displays details for the list of AD members present in that switch.

- Since all E_Ports and EX_Ports are shared across all Admin Domains, they are shown under all Admin Domains.
- Other ports are displayed without any attribute details (with an explanation that they are not part of the your Admin Domain).

A port or device appears in CLI command output or other management tool outputs if any one of the conditions listed in Table 7-8 is met:

Table 7-8 Ports and Devices in CLI Output

For	Condition
(Domain, Port)	• The port is specified in the (domain, port) member list of the Admin Domain. • One or more WWNs specified in the AD member list is attached to the (domain, port).
Device WWN	• The device WWN is specified in the AD WWN member list. • The device WWN is attached to one of the (domain, port) specified in the AD member list.

RASlog and SYSlog output is not filtered based on AD membership.

See the *Fabric OS Command Reference* for more detailed information about command syntax and usage and to understand how existing commands behave in an AD context.

Executing a Command in a Different AD Context

The **ad --exec** option executes a command in a different Admin Domain. The Admin Domain must be one to which you can access. This option creates a new shell with the current user_id, switches to the specified Admin Domain, performs the specified command, and exits out of the shell.

To execute a command in a different Admin Domain context

1. Connect to the switch and log in.
2. Enter the **ad --exec** command, specifying the Admin Domain and the command you want to execute.

   ad --exec *ad_id* "*command*"

   The following example executes the **switchShow** command in the AD7 context.

   ```
   sw5:AD255:admin> ad --exec 7 "switchshow"
   ```

Displaying an Admin Domain Configuration

The **ad --show** option displays the membership information and zone database information of the specified Admin Domain. When you perform the **show** option in:

- AD255, if you do not specify the AD_name or number, all information about all existing Admin Domains is displayed.
- AD0-AD254 contexts, the membership of the current Admin Domain is displayed.
- AD0, the device and switch list members are categorized into implicit and explicit member lists.

To show an Admin Domain

1. Connect to the switch and log in as any user type.
2. Enter the **ad --show** command.

   ```
   ad --show
   ```

   If you are in the AD0 context, you can use the **-i** option to display the implicit membership list of AD0; otherwise, only the explicit membership list is displayed.

   ```
   ad --show -i
   ```

   If you are in the AD255 context, all Admin Domain configuration from the transaction buffer, defined configuration, and effective configuration is displayed, unless you use the **-m** option:

   ```
   ad --show ad_id -m mode
   ```

   where *ad_id* is the Admin Domain for which you want to display information and *mode* is one of the following:

   - 0 to display the Admin Domain configuration in the current transaction buffer.
   - 1 to display the Admin Domain configuration stored in the persistent memory (defined configuration).
   - 2 to display the currently enforced Admin Domain configuration (effective configuration).

   The following example displays membership information about AD1.

```
sw5:AD1:admin> ad --show
Current AD Number: 1  AD Name: TheSwitches

Effective configuration:
------------------------

AD Number: 1 AD Name:   TheSwitches     State: Active

        Switch WWN members:             50:06:06:99:00:2a:e9:01;
                                        50:00:51:e0:23:36:f9:01;
                                        50:06:06:98:05:be:99:01;
```

Switching to a Different Admin Domain Context

The **ad --select** option is used to switch between different Admin Domain contexts. This option creates a new shell with a new Admin Domain context. If the corresponding Admin Domain is not yet activated, the **select** option:

- Succeeds for the physical fabric administrator. (This allows the root administrator to switch to an Admin Domain and set up zones before activating the AD).
- Fails for all other users

To switch to a new Admin Domain context

1. Connect to the switch and log in as any user type.
2. Enter the **ad --select** command and the Admin Domain you want to switch to.
3. To leave the new Admin Domain context, exit from the shell.

   You cannot switch to another Admin Domain context from within the shell created by **ad --select**. You must first exit the shell and then issue the **ad --select** command again.

   The following example switches to the AD12 context. Note that the prompt changes to display the Admin Domain.

```
sw5:admin> ad --select 12
sw5:AD12:admin>
```

Performing Zone Validation

If you are working with zones, you should be aware that there is an Admin Domain impact. Zone objects can be part of an Admin Domain. You can use the **zone --validate** command to list all zone members that are not part of the current zone enforcement table. A member might not be part of the zone enforcement table because:

- The device is offline.
- The device is online, but is connected to an AD-unaware switch.
- The device is online but is not part of the current Admin Domain.

For more information about the **zone** command and its use with Admin Domains, see the *Fabric OS Command Reference*.

Admin Domain Interactions

The administrative domain feature provides interaction with other Fabric OS features and across third-party applications. You can manage Admin Domains with Web Tools applications and with Fabric Manager. If the current Admin Domain owns the switch, you can perform Fabric Watch operations.

Admin Domain interactions do not extend to user session tunneling across switches. A user logged into a switch can control only the local switch ports as specified in the Admin Domain.

When the fabric is in secure mode, the following applies:

- There is no support for ACL configuration under each Administrative Domain.
- ACL configuration commands are allowed only in AD0 and AD255. None of the policy configurations are validated with AD membership.
- You cannot use Admin Domains and Secure Fabric OS in combination. The Secure Fabric OS environment does not support Admin Domains:
  - If Secure Fabric OS is active, you cannot configure Admin Domains.
  - If Admin Domains are configured, you cannot use Secure Fabric OS.

Table 7-9 lists some of the Fabric OS features and considerations that apply when using Admin Domains.

Table 7-9 Admin Domain Interaction with Fabric OS Features

Fabric OS Feature	Admin Domain Interaction
ACLs	If no user-defined Admin Domains exist, you can run ACL configuration commands in only AD0 and AD255. If any user-defined Admin Domains exist, you can run ACL configuration commands only in AD255. You *cannot* use ACL configuration commands or validate ACL policy configurations against AD membership under each Admin Domain.
Advanced Performance Monitoring (APM)	All APM-related filter setup and statistics viewing is allowed only if the local switch is part of the current Admin Domain.
Fabric Watch	Fabric Watch configuration operations are allowed only if the local switch is part of the current Admin Domain.
FCR	You can create LSAN zones as a physical fabric administrator or as an individual AD administrator. The LSAN zone can be part of the root zone database or the AD zone database. • FCR collects the LSAN zones from all ADs. If both edge fabrics have matching LSAN zones and both devices are online, FCR triggers a device import. • LSAN zone enforcement in the local fabric occurs only if the AD member list contains both of the devices (local and imported device) specified in the LSAN zone. To support legacy applications, WWNs are reported based on the AD context using NAA=5. As a result, you cannot use the NAA=5 field alone in the WWN to detect an FC Router.
FDMI	FDMI operations are allowed only in AD0 and AD255.
FICON	Admin Domains support FICON. However, you must perform additional steps because FICON management (CUP) requires additional physical control of the ports. You must set up the switch as a physical member of the FICON AD. DCC and SCC policies are supported only in AD0 and AD255, since ACL configurations are supported only in AD0 and AD255.
iSCSI	iSCSI operations are supported only in AD0.
Management applications	Management interfaces that access the fabric without a user's credentials continue to get the physical fabric view. Examples include: SNMPV1, Web Tools, Fabric Manager; http access, unzoned management server query, FAL in-band CT requests from FAL Proxy to FAL Target, and FC-CT based management applications (such as Tivoli). Access from applications or hosts using Management Server calls can be controlled using the Management Server ACL support provided by the **msConfigure** command. Note that this is a switch-specific setting and not a fabric-wide setting.

Table 7-9 Admin Domain Interaction with Fabric OS Features (Continued)

Fabric OS Feature	Admin Domain Interaction
Port-Swapping and PID formats	Admin Domain port members are specified in (domain, port) format. Based on the PID format, a (domain, port) member indicates a slot/port in the switch. The (domain, port) member is effectively a member of that AD. If the PID format changes, all (domain, port) AD members with a port index of less than 128 are automatically converted to meet the new PID format. Port swapping has no effect on AD support as port swapping swaps only the area numbers of two ports and Admin Domains are specified using (domain, port) member. For detailed information about configuring the PID format, see Appendix A, "Configuring the PID Format".
RSCN	Admin Domains do not introduce any RSCN changes to devices or hosts.

Refer to the *Fabric OS Command Reference* for detailed information about the commands. The following sections describe Admin Domain interactions with zones, zone databases, and LSAN zones.

Admin Domains, Zones, and Zone Databases

Each Admin Domain has its own zone database, with both defined and effective zone configurations and all related zone objects (zones, zone aliases, and zone members). Within an Admin Domain, you can configure zoning only with the devices that are present in that Admin Domain.

With a hierarchical zoning model, the name space for each Admin Domain and the root zones are separate; configurations are supported with the same zone object name appearing in the root zone database and different ADs (for example, the same zoneset name in AD1 and AD2).

Zoning operations ignore any resources not in the Admin Domain, even if they are specified in the zone. The behavior functions similarly to specifying offline devices in a zone. All zones from each Admin Domain zoneset are enforced. The enforcement policy encompasses zones in the effective zoneset of the root zone database and the effective zonesets of each AD.

Note

AD zone databases do not have an enforced size limit. The zone database size is calculated by the upper limit of the AD membership definition and the sum of all the zone databases for each AD.

Admin Domains support **defzone** mode of **noaccess** alone. Before configuring any Admin Domain, you must set the **defzone** to **noaccess** mode. Admin Domains without effective zone configurations are presented with **no access**. See "Activating Default Zones" on page 20-15 for more information.

Admin Domains introduce two types of zone database nomenclature and behavior:

- **Root zone database**—If you do not use Admin Domains, you will have only one zone database. This legacy zone database is known as the *root zone database*. If you create Admin Domains, you will have the root zone database, which is owned by AD0, and other zone databases, one for each user-defined Admin Domain.
  - During the zone update process, only the root zone database is sent to AD-unaware switches.
  - AD-level zone information is merged with the root zone configuration and enforced.

If the administrative domain feature is not active (AD1–AD254 are not configured and no explicit members are added to AD0), AD0 supports both **defzone allaccess** and **noaccess** modes.

- **Zone databases—**The Admin Domains each have separate zone databases and zone transaction buffers. You can concurrently edit the separate zone databases. The AD zone database also has the following characteristics:
  - Each Admin Domain (AD1 through AD254) has its own zone definitions. These zone definitions include defined and effective zonesets and all related zone objects including zones, zone aliases, and zone members. For example, you can define a zone name of test_z1 in more than one Admin Domain.
  - Each zone database has its own namespace.
  - There is no zone database linked to the physical fabric (AD255) and no support for zone database updates. In the physical fabric context (AD255), you can only view the complete hierarchical zone database, which is made up of the zone databases in AD0 through A254.
  - With AD support, zoning updates are supported selectively at each AD level. For example, a zone change in AD1 results in an update request only for the AD1 zone database.

Admin Domains and LSAN Zones

LSANs under each Admin Domain are collated into a single name space and sent out to FCR phantom domains using the following format:

*<original_LSAN_name>*_AD*<AD_num>*

For example, a zone with name *lsan_for_linux_farm* in AD5 is internally converted to *lsan_for_linux_farm_AD005*.

LSAN zone names in AD0 are never converted for backward compatibility reasons.

The auto-converted LSAN zone names might collide with LSAN zone names in AD0 (for example, in the above example, if AD0 contains *lsan_for_linux_farm_AD005*, this would cause a name collision). Fabric OS does not detect or report such name clash.

LSAN zone names greater than 57 characters are not converted or sent to the FCR phantom domain.

Configuration Upload and Download in an AD Context

The behavior of **configUpload** and **configDownload** varies depending on the AD context and whether the switch is a member of the current Admin Domain. In the AD context, these commands include only the zone configuration of the current Admin Domain. If the switch is a member of the Admin Domain, all switch configuration parameters are saved and the zone database for that Admin Domain is also saved. Table 7-10 lists the sections in the configuration file and the Admin Domain contexts in which you can upload and download these sections.

The **configDefault** command does not clear zone or Admin Domain database information. This command is allowed only if the switch is a member of the current Admin Domain.

See Chapter 6, "Maintaining Configurations" for additional information.

Table 7-10 Configuration Upload and Download Scenarios in an AD Context

AD Contexts	Configuration File Sections					
	iSCSI	ACL	Secure Fabric OS	Zone	AD Headers	Switch configuration and other parameters
AD255						
With ADs	Yes	Yes	No	Yes[a]	Yes	Yes
Without ADs	Yes	Yes	Yes	Yes[a]	Yes	Yes
AD0						
With ADs and switch membership	Yes	No	No	Yes[b]	No	Yes
With ADs and without switch membership	Yes	No	No	Yes[b]	No	No
Without ADs	Yes	Yes	Yes	Yes[b]	No	Yes
AD1 – AD254						
With switch membership	No	No	No	Yes[c]	No	Yes
Without switch membership	No	No	No	Yes[c]	No	No

a. Zone databases for AD0 through AD254.
b. Only zone database for AD0.
c. Only zone database for current AD.

Installing and Maintaining Firmware

Chapter 8

This chapter contains procedures for installing and maintaining firmware. Fabric OS v5.2.0 provides nondisruptive firmware installation.

- "About the Firmware Download Process" on page 8-2
- "Preparing for Firmware Downloads" on page 8-3
- "Performing Firmware Download on Switches" on page 8-7
- "Performing Firmware Download on Directors" on page 8-9
- "Testing and Restoring Firmware on Switches" on page 8-14
- "Testing and Restoring Firmware on Directors" on page 8-15
- "Validating Firmware Download" on page 8-19
- "Troubleshooting Firmware Download" on page 8-20

In this chapter, you will see references to specific types of blades inserted into the SilkWorm 48000:

- Port blades contain Fibre Channel ports.
- FC blades contain only Fibre Channel ports: SilkWorm FC4-16/32/48.
- AP blades are "intelligent" blades: SilkWorm FR4-18i and SilkWorm FC4-16IP.
- CP blades have a control processor (CP) used to control the entire switch; they can be inserted only into slots 5 and 6.

About the Firmware Download Process

The **firmwareDownload** command downloads uncompressed Fabric OS firmware from an FTP server to the switch's nonvolatile storage area.

In the SilkWorm 24000 and 48000 directors, this command (without any options flags) by default downloads the firmware image to a standby CP, if there is one, to prevent disruption to application services. This operation depends on high-availability (HA) support. If HA is not available, you can upgrade the CPs one at a time, using the **-s** option.

Caution
To ensure nondisruptive downloads on switches (versus directors), ensure that the firmware is download is completed successfully on each switch in a serial fashion. Enter the **firmwareDownloadStatus** command on the switch before moving on to the next switch.

SilkWorm fixed-port models and each CP blade of the SilkWorm 24000 and 48000 models have two partitions of nonvolatile storage areas (a primary and a secondary) to store two firmware images. The **firmwareDownload** command always loads the new image into the secondary partition and swaps the secondary partition to be the primary. It then reboots the switch/CP and activates the new firmware. Finally, **firmwareCommit** starts after each CP has one partition upgraded (on fixed ports, after the reboot), to copy the new image to the other partition.

If you are using a SilkWorm 48000 with one or more AP blades: The Fabric OS automatically detects mismatches between the active CP firmware and the blade's firmware. By the end of the firmware download process the active CP and the blade will be running the same version of firmware.

Upgrading and Downgrading Firmware

"Upgrading" means installing a newer version of firmware; "downgrading" means installing an older version of firmware.

In most cases, you will be upgrading firmware; that is, installing a newer firmware version than the one you are currently running. However, some circumstances might require installing an older version; that is, downgrading the firmware. The procedures in this section assume that you are upgrading firmware, but they work for downgrading as well, provided the old and new firmware versions are compatible. Always reference the latest release notes for updates that may exist regarding downgrades under particular circumstances.

Using the CLI (or Web Tools), you can upgrade the firmware on one switch at a time. You can use the optionally licensed Brocade Fabric Manager application to upgrade firmware simultaneously on multiple switches. For more details on Fabric Manager and other licensed software tools, go to the Brocade Web site at *http://www.brocade.com*.

Effects of Firmware Changes on Accounts and Passwords

The following table describes what happens to accounts and passwords when you replace the switch firmware with a different version.

Table 8-1 Effects of Firmware Changes on Accounts and Passwords

Change	First Time	Subsequent Times (After upgrade, then downgrade, then upgrade)
Upgrading	Default accounts and their passwords are preserved.	User-defined and default accounts and their passwords are preserved.
Downgrading	User-defined accounts are no longer valid. Default accounts and their passwords are preserved. If a default account was disabled, it is reenabled after the downgrade.	User-defined and default accounts and their passwords are preserved, including accounts added after the first upgrade.
Upgrading to v3.2.0	(You might upgrade a switch in the fabric as part of "Checking Connected Switches" on page 8-5.) Earlier versions allowed you to change the default account names. You cannot add user-defined accounts until you change the names back to default with the **passwdDefault** command.	

For details on administrative domains and firmware download, see Chapter 7, "Managing Administrative Domains". For more details on older releases of Fabric OS, see "Understanding Legacy Password Behavior" on page C-1.

For details about downgrading and restoring firmware, see "Testing and Restoring Firmware on Directors" on page 8-15.

Considerations for FICON CUP Environments

To prevent channel errors during nondisruptive firmware installation, the switch CUP port must be taken offline from all host systems.

Preparing for Firmware Downloads

Before executing a firmware download, it is recommended that you perform the tasks mentioned in this section. In the unlikely event of a failure or timeout, the preparation tasks that are described in this section will enable you to provide Brocade Support all the information required to perform advanced troubleshooting.

To provide specific information about why a firmware downgrade might fail, starting in Fabric OS v5.2.0, the following exception cases are checked. If you're attempting to downgrade to Fabric OS v5.1.x or lower (or in some cases v4.4.x or lower) and the check finds that one of these exception cases is true, firmware download will fail and an error message will be displayed.

It is recommended that you perform a **configUpload** to back up the current configuration before you download firmware to a switch. See "Backing Up a Configuration" on page 6-2 for details.

To prepare for a firmware download

1. Read the release notes for the new firmware to find out if there are any additional issues related to firmware download.
2. Establish a telnet session and log in as admin. Enter the **firmwareShow** command to verify the current version of Fabric OS. Brocade does not support upgrades from more than two previous releases. For example upgrading from Fabric OS v4.4.0 to v5.1.x is supported, but upgrading from Fabric OS v4.4.0 or a previous release directly to v5.2.0 is not.

   In other words, upgrading a switch from Fabric OS v4.4.0 or a previous release to v5.2.0 is a two-step process—first upgrade to v5.0.0 or v5.0.1 and then upgrade to v5.2.0. The major upgrade versions are:

   - Fabric OS v4.0.0, v4.1.0, v4.2.0, and v4.4.0
   - Fabric OS v5.0.1, v5.1.0, and v5.2.0

   Caution
   Newer Fabric OS versions (v4.4.0 and above) can support large zone databases. However, exercise caution when downgrading to v4.2.0 and earlier versions that do not support large zone databases. If the zone database is too large, then you may be required to disable all ports in earlier versions of firmware to avoid data corruption; this also requires a manual reactivation of all ports.

   Fabric OS v4.4.0 through v5.1.0 support zone databases up to 256 KB in size and then up to 1 MB in v5.2.0. If you upgrade to Fabric OS v5.2.0 from v4.4.x or later and then want to downgrade, you must reduce the size of the zone database to 256 KB or less.
3. Upon initial setup of a factory-delivered switch, make sure that all IP address have been set and the switch has been rebooted prior to running a firmware download. Establish a telnet session and log in to the SilkWorm switch.
4. (Optional) For enhanced support, connect the switch with a serial console cable to a computer. Ensure that all serial consoles (both CPs for bladed products) and all telnet sessions are logged and included with any trouble reports.
5. Establish a telnet session and log in to the switch as admin. Issue the **saveCore** command to remove all extra core files prior to executing the firmware download. This helps free up some disk space. If you are upgrading a director, log in to both the active and standby CPs and enter the **saveCore** command on each CP.

   Note
   If you do not know the CP address, use the **ipAddrShow** command to view a list of all CP IP addresses associated with the switch.
6. (Optional) Enter the **supportSave** command to capture a snapshot of the configuration and provide baseline information in case there is a need to troubleshoot or seek advanced support. Make sure that you enter this command on the standby CP as well.

After you run **supportSave**, locate the compact flash usage information in the **supportSave** output files named *XXX_SUPPORTSHOW* (where *XXX* is a switch, date, and time prefix). Open the *XXX_SUPPORTSHOW* file in any text editor and search for the following information:

```
/bin/df:
Filesystem           1k-blocks      Used Available Use% Mounted on
/dev/root               120112     61128     58984  51% /
/dev/hda2               120128     72356     47772  60% /mnt
```

Verify that the compact flash usage is not above 90%. If the compact flash usage is above 90%, contact your switch service provider.

Note

If using Fabric OS v4.2.0 or earlier, enter the **supportShow** command and verify the above compact flash information by searching the output of the **supportShow** command.

7. (Optional) Enter the **errClear** command to erase all existing messages in addition to internal messages.

Checking Connected Switches

If the switch to be upgraded is running v4.1.0 firmware or later, it is recommended that all switches directly connected to it be running versions no earlier than v2.6.1, v3.1.0, or v4.1.0. If some connected switches are running older firmware, upgrade them to *at least* the earliest recommended version (shown in Table 8-2) before upgrading firmware on the switch.

Note

Go to *http://www.brocade.com* to view end-of-life policies for Brocade products. Navigate to the Services and Support tab, then select End of Life Information. End-of-life products are not supported.

Table 8-2 Recommended Firmware

SilkWorm Model	Earliest Recommended Fabric OS Version
200E	v5.0.1
2000 series	v2.6.1
3200, 3600, 3800	v3.1.0
3014, 3016, 3250, 3850	v4.2.0
3900	v4.1.0
4012, 4020	v5.0.0
4100	v4.4.0
4900	v5.1.0
7500	v5.1.0

Table 8-2 Recommended Firmware (Continued)

SilkWorm Model	Earliest Recommended Fabric OS Version
24000	v4.2.0
48000	v5.0.1
48000 with FR4-18i blade(s)	v5.1.0
48000 with FC4-48 and/or FC4-16IP blade(s)	v5.2.0

During code activation on SilkWorm 3250, 3850, or 3900 running Fabric OS v4.1.0 or later, data continues to flow between hosts and storage devices; however, fabric services are unavailable for a period of approximately 50 to 55 seconds. Possible disruption of the fabric can be minimized by ensuring that switches logically adjacent to these models (directly connected via an ISL) are running at the minimum Fabric OS v2.6.1 or later, v3.1.0 or later, or v4.1.0 or later.

If SilkWorm 3250, 3850, 3900, 4100, 4900 or 7500 models are adjacent and you start firmware downloads on them at same time, there might be traffic disruption.

To determine whether you need to upgrade connected switches before upgrading the switch, use the following procedure on each connected switch to display firmware information and build dates.

1. Connect to the switch and log in as admin.
2. Enter the **version** command.

   The following information is displayed:

   - **Kernel** displays the version of switch kernel operating system.
   - **Fabric OS** displays the version of switch Fabric OS.
   - **Made on** displays the build date of firmware running in switch.
   - **Flash** displays the install date of firmware stored in nonvolatile memory.
   - **BootProm** displays the version of the firmware stored in the boot PROM.

Obtaining and Decompressing Firmware

Firmware upgrades are available for customers with support service contracts and partners on the Brocade Web site at *http://www.brocade.com*.

At the Brocade Web site, click *Brocade Connect* and follow the instructions to register and download firmware. Partners with authorized accounts can use the *Brocade Partner Network*.

You must decompress the firmware (using the UNIX tar or gzp command for all .gzp files, or a Windows unzip program for all .zip files) before you can use the **firmwareDownload** command to update the firmware on your equipment.

When you unpack the downloaded firmware it expands into a directory that is named according to the version of Fabric OS it contains. For example, if you download and unzip Fabric OS v5.2.0.zip, it expands into a directory called v5.2.0. When you issue the **firmwareDownload** command, there is an automatic search for the correct package file type associated with the switch. Specify only the path to the v5.2.0 directory.

Performing Firmware Download on Switches

SilkWorm 200E, 3250, 3850, 3900, 4100, 4900, and 7500 switches also maintain primary and secondary partitions for firmware. The **firmwareDownload** command defaults to an autocommit option that automatically copies the firmware from one partition to the other.

You should not override autocommit under normal circumstances; use the default. See "Testing and Restoring Firmware on Directors" on page 8-15 for details about overriding the autocommit option.

Overview of the Firmware Download Process on Switches

The following section describes the default behavior after you enter the **firmwareDownload** command (without options) on SilkWorm 200E, 3250, 3850, 3900, 4100, 4900, and 7500 models:

- The Fabric OS downloads the firmware to the secondary partition.
- The system performs a high-availability reboot (**haReboot**). After the **haReboot**, the former secondary partition is the primary partition.
- The system replicates the firmware from the primary to the secondary partition.

SilkWorm 200E, 3250, 3850, 3900, 4100, 4900 and 7500 Firmware Download

The upgrade process first downloads and then commits the firmware to the switch. While the upgrade is proceeding, you can start another telnet session on the switch and use the **firmwareDownloadStatus** command to observe the upgrade progress if you wish.

Note

After you start the process, do not enter any disruptive commands (such as **reboot**) that will interrupt the process. The entire firmware download and commit process takes approximately 17 minutes. If there is a problem, wait for the timeout (30 minutes for network problems). Disrupting the process can render the switch inoperable and require you to seek help from your switch service provider.

Do not disconnect the switch from power during the process, because the switch could become inoperable upon reboot.

To upgrade firmware for SilkWorm 200E, 3250, 3850, 3900, 4100, 4900, and 7500 switches

1. Verify that the FTP service is running on the host server and that you have a user ID on that server.
2. Obtain the firmware file from the Brocade Web site at *http://www.brocade.com* and store the file on the FTP server. Verify that the FTP service is running and unpack the compressed files preserving directory structures.
3. Use the **firmwareShow** command to check the current firmware version on connected switches. Upgrade their firmware if necessary before proceeding with upgrading this switch.

   See "Checking Connected Switches" on page 8-5 for details.

4. Connect to the switch and log in as admin.
5. Enter the **firmwareDownload** command.
6. At the “Do you want to continue [y/n]” prompt, enter **y**.
7. Respond to the prompts as follows:

Server Name or IP Address	Enter the name or IP address of the FTP server where the firmware file is stored; for example, 192.1.2.3. You can enter a server name if DNS is enabled.
User name	Enter the user name of your account on the server; for example, “JaneDoe”.
File name	**v5.2.0 or higher**: Specify the full path name of the firmware directory, for example, */pub/v5.2.0*. **v5.1.0 or lower**: Specify the full path name of the firmware directory, appended by *release.plist*; for example, */pub/v5.1.0/release.plist*. **Note**: For v4.x and v5.x switches only, do not attempt to locate the *release.plist* file in the top level directory; there is a *release.plist* file for each platform, and the correct one is automatically selected.
Password	Enter your account password for the server. Note that this is a required field even if you are logged in as an anonymous user; in such cases, the value may be ignored by the FTP service.

After the firmware is downloaded, the switch performs an HA reboot and starts the firmware commit.

8. After the HA reboot, connect to the switch and log in again as admin.
9. If you want snapshots of the upgrade progress, using a separate telnet session, enter the **firmwareDownloadStatus** command to monitor the firmware download.
10. After the firmware commit is completed, enter the **firmwareShow** command to display the firmware level of both partitions. It takes several minutes to complete the commit operation.

```
switch:admin> firmwaredownload
Server Name or IP Address: 10.1.2.3
FTP User Name: JaneDoe
File Name: /pub/v5.2.0
FTP Password: xxxx
Verifying the input parameters ...
Checking system settings for firmwaredownload...
You can run firmwaredownloadstatus to get the status
of this command.

This command will cause a warm/non-disruptive boot on the switch,
but will require that existing telnet, secure telnet or SSH sessions
be restarted.

Do you want to continue [Y] y
```

Performing Firmware Download on Directors

You can download firmware to SilkWorm 24000 and 48000 directors without disrupting the overall fabric if the two CP blades are installed and fully synchronized. Use the **haShow** command to confirm synchronization. If only one CP blade is powered on, the switch must be rebooted to activate firmware, which is disruptive to the overall fabric.

During the upgrade process, the director fails over to its standby CP blade and the IP addresses for the two logical switches move to that CP blade's Ethernet port. This might cause informational ARP address reassignment messages to appear on other switches in the fabric. This is normal behavior, because the association between the IP addresses and MAC addresses has changed.

Caution
To successfully download firmware you must have an active Ethernet connection on *each* of the CPs.

If your fabric is set to the extended edge PID format and you want to downgrade to an older Fabric OS version that does not support extended edge, you *must* change the PID to a supported format. See Appendix A, “Configuring the PID Format” for more information about PID formats.

Use the **haShow** command to verify that the CPs are synchronized prior to running a firmware download. If the CPs are not in sync, run the **haSyncStart** command. If the problem persists, review “Troubleshooting Firmware Download” on page 8-20. If the troubleshooting information fails to help resolve the issue, contact your switch service provider.

Overview of the Firmware Download Process on Directors

The following summary describes the default behavior of the **firmwareDownload** command (without options) on SilkWorm 24000 and 48000 directors.

After you enter the **firmwareDownload** command on the active CP blade:

- The standby CP blade downloads firmware.
- The standby CP blade reboots and comes up with the new Fabric OS.
- The active CP blade synchronizes its state with the standby CP blade.
- The active CP blade forces a failover and reboots to become the standby CP blade.
- The *new* standby CP blade (the active CP blade before the failover) downloads firmware.
- The *new* standby CP blade reboots and comes up with the new Fabric OS.
- The new active CP blade synchronizes its state with the new standby CP blade.

- The **firmwareCommit** command runs automatically on both CP blades.

Note

After you start the process, do not enter any disruptive commands (such as reboot) that will interrupt the process. The entire firmware download and commit process takes approximately 15 minutes. If there is a problem, wait for the timeout (30 minutes for network problems). Disrupting the process can render the switch inoperable and require you to seek help from your switch vendor.

Do not disconnect the switch from power during the process, because the switch could become inoperable upon reboot.

SilkWorm 24000 and 48000 Firmware Download Procedure

There is one logical switch address for a SilkWorm 48000, and up to two logical switch addresses for the SilkWorm 24000, but either can be used on the SilkWorm 24000 to effect a firmwaredownload (either logical switch).

Note

By default, the **firmwareDownload** command automatically upgrades both the active CP blade and the standby CP blade; it automatically upgrades all AP blades in the SilkWorm 48000.

To upgrade firmware on SilkWorm 24000 and 48000 directors (including blades)

1. Verify that the FTP service is running on the host server and that you have a user ID on that server.
2. Obtain the firmware file from the Brocade Web site at *http://www.brocade.com* and store the file on the FTP server. Verify that the FTP service is running and decompress the compressed files preserving directory structures.
3. Use the **firmwareShow** command to check the current firmware version on connected switches. Upgrade the firmware, if necessary, before proceeding with upgrading this switch.

Caution

SilkWorm 48000 with an FR4-18i blade: If you are running v5.1.0 firmware, then you cannot downgrade to earlier versions without removing the blade(s).

SilkWorm 48000 with anFC4-48 or FC4-16IP blade: If you are running v5.2.0 firmware, then you cannot downgrade to earlier versions without removing the blade(s).

Do not remove blades until the port configuration, for example, EX_Ports, are removed first. The **firmwareDownload** command will indicate when blades are save to remove.

See “Checking Connected Switches” on page 8-5.

4. Using a telnet session, connect to the switch and log in as admin.

5. Enter the **haShow** command to confirm that the two CP blades are synchronized. In the following example, the active CP blade is CP1 and the standby CP blade is CP0:

```
switch:admin> hashow
Local CP (Slot 5, CP0): Active, Warm Recovered
Remote CP (Slot 6, CP1): Standby, Healthy
HA enabled, Heartbeat Up, HA State synchronized
switch:admin>
```

   CP blades must be synchronized and running Fabric OS v4.2.0 or later to provide a nondisruptive download. If the two CP blades are not synchronized, enter the **haSyncStart** command to synchronize them. If the CPs still are not synchronized, contact your switch service provider.

6. Enter the **firmwareDownload** command.
7. At the "Do you want to continue [y/n]" prompt, enter **y.**
8. Respond to the prompts as follows:

Server Name or IP Address	Enter the name or IP address of the server where the firmware file is stored: for example, 192.1.2.3. You can enter a server name if DNS is enabled.
User name	Enter the user name of your account on the server: for example, *JaneDoe*.
File name	**v5.2.0 or higher:** Specify the full path name of the firmware directory, for example, */pub/v5.2.0.* **v5.1.0 or lower:** Specify the full path name of the firmware directory, appended by *release.plist*; for example, */pub/v5.1.0/release.plist*. **Note**: For v4.x and v5.x switches only, do not attempt to locate the *release.plist* file in the top level directory; there is a *release.plist* file for each platform, and the correct one is automatically selected.
Password	Enter your account password for the server. Note that this is a required field even if you are logged in as an anonymous user; in such cases, the value may be ignored by the FTP service

   The firmware is downloaded to one CP blade at a time, beginning with the standby CP blade. During the process, the active CP blade is failed over. After the firmware is downloaded, a firmware commit starts on both CP blades. The entire firmware download and commit process takes approximately 35 minutes.

   If there is an AP blade present: At the point of the failover an *autoleveling* process is activated. autoleveling is triggered when the active CP detects a blade that contains a different version of the firmware, regardless of which version is older. Autoleveling downloads firmware to the AP blade, swaps partitions, reboots the blade, and copies the new firmware from the primary partition to the secondary partition. If you have multiple AP blades, they are updated simultaneously; however, the downloads can occur at different rates.

Autoleveling takes place in parallel with the firmware download being performed on the CPs, but does not impact performance. Fibre Channel traffic is not disrupted during autoleveling, but GbE traffic on AP blades may be affected.

```
switch:admin> firmwaredownload
Server Name or IP Address: 10.1.2.3
FTP User Name: JaneDoe
File Name: /pub/v5.2.0
FTP Password: xxxx
Verifying the input parameters ...
Checking system settings for firmwaredownload...
The following AP blades are installed in the system.

Slot Name      Versions                          Traffic Disrupted
-------------------------------------------------------------------
  3  FC4-16IP  v5.2.0_main_bld47                       GigE
  4  FR4-18i   v5.2.0_main_bld47                       None
 10  FR4-18i   v5.2.0_main_bld47                       None

This command will update the firmware on both CPs and all AP blade(s)
above.  If you want to update firmware on a single CP only,
please use -s option.

You can run firmwaredownloadstatus to get the status of this
command.

This command will cause a warm/non-disruptive boot on the active CP,
but will require that existing telnet, secure telnet or SSH sessions
be restarted.

Do you want to continue [Y]: y
```

9. Optionally, after the failover, connect to the switch, log in again as admin. Using a separate telnet session, enter the **firmwareDownloadStatus** command to monitor the firmware download status.

```
switch:admin> firmwaredownloadstatus
[1]: Fri Sep 22 09:45:15 2006
Slot 5 (CP0, active): Firmware is being downloaded to standby CP. This step may
take up to 30 minutes.

[2]: Fri Sep 22 09:51:21 2006
Slot 5 (CP0, active): Firmware has been downloaded successfully to Standby CP.

[3]: Fri Sep 22 09:51:28 2006
Slot 5 (CP0, active): Standby CP is going to reboot with new firmware.

[4]: Fri Sep 22 09:53:00 2006
Slot 5 (CP0, active): Standby CP booted successfully with new firmware.

[5]: Fri Sep 22 09:53:11 2006
Slot 10 (FC4-16IP): Firmware is being downloaded to the blade. This step may take
up to 10 minutes.

[6]: Fri Sep 22 09:53:54 2006
Slot 6 (CP1, active): Forced failover succeeded. New Active CP is running new
firmware

[7]: Fri Sep 22 09:54:01 2006
Slot 6 (CP1, active): Firmware is being download to standby CP. This step may
take up to 30 minutes.

[8]: Fri Sep 22 09:54:59 2006
Slot 10 (FC4-16IP): Firmware has been downloaded successfully. Blade is rebooting
with the new firmware.

[9]: Fri Sep 22 09:55:36 2006
Slot 10 (FC4-16IP): Firmware commit has started on the blade. This may take up to
10 minutes.

[10]: Fri Sep 22 09:56:19 2006
Slot 10 (FC4-16IP): The commit operation has completed successfully.

[11]: Fri Sep 22 10:00:30 2006
Slot 6 (CP1, active): Firmware has been downloaded successfully on Standby CP.

[12]: Fri Sep 22 10:00:38 2006
Slot 6 (CP1, active): Standby CP reboots.

[13]: Fri Sep 22 10:02:16 2006
Slot 6 (CP1, active): Standby CP booted successfully with new firmware.

[14]: Fri Sep 22 10:02:20 2006
Slot 6 (CP1, active): Firmware commit operation has started on both active and
standby CPs.

[15]: Fri Sep 22 10:06:08 2006
Slot 6 (CP1, active): Firmware commit operation has completed successfully on
both CPs.

[16]: Fri Sep 22 10:06:08 2006
Slot 6 (CP1, active): Firmwaredownload command has completed successfully. Use
firmwareshow to verify the firmware versio
ns.

switch:admin>
```

10. Enter the **firmwareShow** command to display the new firmware versions.:

```
switch:admin> firmwareshow
Slot Name      Primary/Secondary Versions        Status
-----------------------------------------------------------------
  3  FC4-16IP  v5.2.0
               v5.2.0
  4  FR4-18i   v5.2.0
               v5.2.0
  5  CP0       v5.2.0                           ACTIVE *
               v5.2.0
  6  CP1       v5.2.0                           STANDBY
               v5.2.0
 10  FR4-18i   v5.2.0
               v5.2.0
switch:admin>
```

Testing and Restoring Firmware on Switches

Typically, users downgrade firmware after briefly “test driving” a newer (or older) version and then restore the original version of the firmware. Testing a new version of firmware in this manner ensures that you do not replace existing firmware because the test drive version occupies only one partition on the switch.

Caution
When you test drive new firmware make sure you have disabled all features that are not supported by the original firmware before restoring to the original version.

1. Prepare

a. Start a telnet session to the Brocade logical switch IP address.

b. Enter the **firmwareShow** command to view the current firmware.

2. Update firmware on the switch

a. Enter the **firmwareDownload -s** command and respond to the prompts as follows:

```
switch:admin> firmwaredownload -s
Server Name or IP Address: 10.32.220.100
FTP User Name: JaneDoe
File Name: /pub/v5.2.0
FTP Password:
Do Auto-Commit after Reboot [Y]: n
Reboot system after download [N]: y
Firmware is being downloaded to the switch. This step may take up to 30
minutes.
Checking system settings for firmwaredownload...
```

The switch will reboot and come up with the new firmware to be tested. Your telnet session will be automatically disconnected.

b. Start a new telnet session, log in as admin, and enter the **firmwareShow** command to confirm that the primary partition of the switch contains the new firmware.

You are now ready to evaluate the new version of firmware.

Note

Stop! If you wish to *restore* the firmware, stop here and skip ahead to step 4; otherwise, continue to step 3 to commit the firmware on the switch, which completes the firmware download operations.

3. Commit the firmware

a. Enter the **firmwareCommit** command to update the secondary partition with new firmware. Note that it takes several minutes to complete the commit operation.

b. Enter the **firmwareShow** command to confirm both partitions on the switch contain the new firmware.

Note

Stop! If you have completed step 3, then you have committed the firmware on the switch and you have completed the firmware download procedure. Step 4 describes how to restore the original firmware, and should be performed after step 2.

4. Restore the firmware

a. Enter the **firmwareRestore** command. The switch will reboot and come up with the original firmware again.

A **firmwareCommit** will automatically begin to copy the original firmware from the primary partition to the secondary partition. At the end of the firmware commit process, both partitions will have the original firmware. Note that it takes several minutes to complete the commit operation.

b. Wait 5 minutes and log in to the switch. Enter the **firmwareShow** command and verify that both partitions on the switch have the original firmware.

Testing and Restoring Firmware on Directors

This procedure enables you to perform a firmware download on each CP and verify that the procedure was successful before committing to the new firmware. The old firmware is saved in the secondary partition of each CP until you enter the **firmwareCommit** command. If you decide to back out of the installation prior to the **firmwareCommit** you can enter the **firmwareRestore** command to restore the older firmware version.

Typically, users downgrade firmware after briefly “test driving” a newer (or older) version and then restore the original version of the firmware. Testing a new version of firmware in this manner ensures that you do not replace the existing firmware because the test drive version uses a partition on each CP.

Note

Brocade recommends that under normal operating conditions you maintain the same firmware version on both CPs, and on both partitions of each CP. This procedure enables you to test firmware before you commit; however, you should not run mixed firmware levels on CPs as a standard practice.

1. Prepare

a. Start a telnet session to the Brocade logical switch IP address.

b. Enter the **ipAddrShow** command and note the address of CP0 and CP1.

c. Enter the **haShow** command and note which CP is active and which CP is standby. Verify that both CPs are in sync. CP blades must be synchronized and running Fabric OS v4.4.0 or later to provide a nondisruptive download. If the two CP blades are not synchronized, enter the **haSyncStart** command to synchronize them. If the CPs still are not synchronized, contact your switch service provider.

d. Enter the **firmwareShow** command and confirm that the current firmware on both partitions on both CPs is as expected.

e. Exit the telnet session.

2. Update firmware on the standby CP

a. Start a telnet session, log in as admin to the standby CP.

b. Enter the **firmwareDownload -s** command and respond to the prompts as follows:

```
switch:admin> firmwaredownload -s
Server Name or IP Address: 10.1.2.3
FTP User Name: JaneDoe
File Name: /pub/v5.2.0
FTP Password:
Do Auto-Commit after Reboot [Y]: n
Reboot system after download [N]: y
Firmware is being downloaded to the switch. This step may take up to 30
minutes.
Checking system settings for firmwaredownload...
```

At this point the firmware should download to the standby CP only and when it has completed the download to that CP, reboot it. The telnet session will be disconnected.

3. Fail over to the standby CP

a. Start a telnet session on the active CP.

b. Enter the **haShow** command to verify that HA synchronization is complete. It will take a minute or two for the standby CP to reboot and synchronize with the active CP.

If the CPs do not achieve synchronization, *stop here*; log in to the standby CP, and enter the **firmwareRestore** command to restore the original firmware.

c. Enter the **firmwareShow** command to confirm that the primary partition of the standby CP contains the new firmware.

d. Enter the **haFailover** command. The active CP will reboot and the telnet session will be disconnected.

If an AP blade is present: At the point of the failover an *autoleveling* process is activated. See Step 8 in the "SilkWorm 24000 and 48000 Firmware Download Procedure" for details about autoleveling.

4. Verify the failover

a. Start a telnet session on the active CP (which is the old standby CP).

b. Enter the **haShow** command to verify that the HA synchronization is complete. It will take a minute or two for the standby CP (which is the old active CP) to reboot and synchronize with the active CP.

Note

If the CPs fail to synchronize, you can still proceed because the version being tested is already present on the active CP, and subsequent steps will ensure that the standby CP is updated to the same version as the active CP.

c. The test version of firmware is now running on the active CP (you can confirm this by entering the **firmwareShow** command).

5. Update firmware on the standby CP

a. Start a telnet session on the standby CP (which is the old active CP).

b. Enter the **firmwareDownload -s** command and respond to the prompts as follows:

```
switch:admin> firmwaredownload -s
Server Name or IP Address: 10.1.2.3
FTP User Name: JaneDoe
File Name: /pub/v5.2.0
FTP Password:
Do Auto-Commit after Reboot [Y]: n
Reboot system after download [N]: y
Firmware is being downloaded to the switch. This step may take up to 30
minutes.
Checking system settings for firmwaredownload...
```

At this point the firmware should download to the standby CP only and reboot it. The telnet session will be disconnected.

c. Wait one minute for the standby CP to reboot, and then start a telnet session and log in as admin.

d. Enter the **firmwareShow** command to confirm that *both* primary partitions now have the test drive firmware in place.

You are now ready to evaluate the new version of firmware.

Note

Stop! If you wish to *restore* the firmware, stop here and skip ahead to step 8; otherwise, continue to step 6 to commit the firmware on both CPs, which completes the firmware download.

6. Perform commit on the standby CP

From the telnet session on the standby CP, enter the **firmwareCommit** command to update the secondary partition with new firmware. Note that it takes several minutes to complete the commit operation.

7. Perform commit on active CP

a. From the telnet session on the active CP, enter the **firmwareShow** command and confirm that only the active CP secondary partition contains the old firmware.

b. Enter the **firmwareCommit** command to update the secondary partition with new firmware. Note that it takes several minutes to complete the commit operation.

c. Enter the **firmwareShow** command to confirm both partitions on both CPs contain the new firmware.

d. Enter the **haShow** command to confirm that the HA state is in sync.

Note

Stop! If you have completed step 7, then you have committed the firmware on both CPs and you have completed the firmware download procedure. Steps 8 through 10 describe how to restore the original firmware, and should be performed after step 5.

8. Restore the firmware on the standby CP

In the telnet session for the standby CP, enter the **firmwareRestore** command. The standby CP will reboot and the telnet session will end. Both partitions will be made equal after several minutes.

9. Perform haFailover on the active CP

a. In the telnet session for the active CP, enter the **haShow** command to verify that HA synchronization is complete. It will take a minute or two for the standby CP to reboot and synchronize with the active CP.

b. Enter the **haFailover** command. The active CP will reboot and the telnet session will end. The switch is now running the original firmware.

10. Restore firmware on the “new” standby CP

a. Wait one minute and start a telnet session on the new standby CP, which is the old active CP.

b. Enter the **firmwareRestore** command. The standby CP will reboot and the telnet session will end. Both partitions will be made equal after several minutes.

c. Wait 5 minutes and log in to the switch. Enter the **firmwareShow** command and verify that all partitions have the original firmware.

If an AP blade is present: Blade partitions always contain the same version of the firmware on both partitions (it does not keep two copies). The firmware is stored on the blade’s compact flash card and is always synchronized with the active CP’s firmware. Thus, if you restore the active CP firmware, the blade firmware is automatically downloaded to become consistent with the new CP firmware (the blade firmware is basically restored).

Your system is now restored to the original partitions on both CPs. Make sure that servers using the fabric can access their storage devices.

If you wish to upgrade a director with only one CP in it, follow the procedures in "Testing and Restoring Firmware on Switches" on page 8-14. Note, however, that upgrading a director with only one CP will be disruptive to switch traffic.

Validating Firmware Download

Validate the firmware download by running the following commands: **firmwareShow**, **nsShow**, **nsAllShow**, and **fabricShow**.

Note

When you prepared for the firmware download earlier, you issued either the **supportShow** (4.2.0 or earlier) or **supportSave** (4.4.0 or later) command. While you can issue the command again and compare the output from before and after, it may take as long as 30 minutes for the command to execute. To save time, it is recommended that you use the commands listed below, which are all are subsets of the **supportSave** output.

All of the connected servers, storage, and switches should be present in the output of these commands. If there is a discrepancy, it is possible that a device or switch cannot connect to the fabric and further troubleshooting is necessary.

firmwareShow	Displays the current firmware level on the switch. For SilkWorm directors this command displays the firmware loaded on both partitions (primary and secondary) for both CPs and AP blades. Brocade recommends that you maintain the same firmware level on both partitions of each CP within the SilkWorm director.
nsShow	(Optional) Displays all devices directly connected to the switch that have logged into the Name Server. Make sure the number of attached devices after the firmware download is exactly the same as the number of attached devices prior to the firmware download.
nsAllShow	(Optional) Displays all connected devices to a fabric. Make sure the number of attached devices after the firmware download is exactly the same as the number of attached devices prior to the firmware download.
fabricShow	(Optional) Displays all switches in a fabric. Make sure the number of switches in the fabric after the firmware download is exactly the same as the number of attached devices prior to the firmware download.

Note

You cannot perform a firmware downgrade from Fabric OS v5.2.0 or higher if administrative domains are configured in the fabric. See Chapter 7, "Managing Administrative Domains" for details.

When the primary and secondary CPs in a SilkWorm 48000 are running pre-Fabric OS v5.2.0 and are in HA-Sync, if firmware is downloaded to upgrade only one CP (using the **firmwareDownload –s** option), that CP will run in an AD-unaware mode (AD creation operations will fail and the local switch will appear as an AD-unaware switch in the fabric).

If primary and secondary CPs are running Fabric OS v5.2.0 and if ADs are configured, any attempt to downgrade one or both CPs to a pre-Fabric OS v5.1.0 version will fail.

See "Troubleshooting Firmware Download" on page 8-20 for exception cases and associated error messages generated from the preinstallaton check.

Troubleshooting Firmware Download

Starting in Fabric OS v5.2.0, many problems with troubleshooting firmware download are removed with the introduction of a network diagnostic script and preinstallation check. However, you should follow these good practices for firmware download before you start the procedure:

- Keep all telnet sessions and serial console logs.
- Enter the **supportSave** (for Fabric OS v4.4.0 or higher) or the **supportShow** command (for Fabric OS 4.2 or earlier) *before and after* entering the **firmwareDownload** command.
- If a problem persists, package together all of the information (the telnet session logs and serial console logs, output from the **supportSave** command) for support. Make sure you identify which information was gathered before and after issuing the **firmwareDownload** command.

If the firmware download fails, see the *Fabric OS Message Reference* for details about error messages. See also "Downgrading Firmware from Fabric OS v5.2.0" next.

If a firmware download fails in a director, the **firmwareDownload** command synchronizes the firmware on the two partitions of each CP by starting a firmware commit operation. Wait *at least* 10 minutes for this commit operation to complete before attempting another firmware download.

If the firmware download fails in a director, the CPs may end up with different versions of firmware and be unable to achieve HA synchronization. In such cases, issue the **firmwareDownload -s** command; the single mode (-s) option upgrades the firmware on the standby CP to match the firmware version running on the active CP. Then re-issue the **firmwareDownload** command to download the desired firmware version to both CPs. For example, if CP0 is running v5.0.1 on the primary and secondary partitions, and CP1 is running v4.4.0e on the primary and secondary partition, then synchronize them by issuing the **firmwareDownload -s** command.

See the *Fabric OS Message Reference* for detailed information about .plist-related error messages.

Downgrading Firmware from Fabric OS v5.2.0

Because of significant features added in Fabric OS in v5.2.0, such as administrative domains, constraints for downgrading firmware exist. Attempts to downgrade from v5.2.0 are blocked if any of the following conditions are true:

- Administrative domains are enabled.
- Port mirroring, fastwrite, or tape pipelining is enabled.
- The zone database is greater than 256 KB. To discover the size of the zone database, use the **cfgSize** command.
- The **timeZone** command has been set to new values.
- RADIUS configuration is set to Remote and Local.
- EX_Port trunking is enabled (feature not supported in Fabric OS v5.1.0).

For more information on any of the commands in the Recommended Action section, see the *Fabric OS Command Reference*.

Note

Some of the messages include error codes (as shown in the example below). These error codes are for internal use only and you can disregard them.

Example: Port configuration with EX ports enabled along with trunking for port(s) 63, use the portcfgexport, portcfgvexport, and/or portcfgtrunkport commands to remedy this. Verify blade is ENABLED. (error 3)

Preinstallation Messages

The messages in this section are displayed if an exception case is encountered during firmware download from Fabric OS v5.2.0. The example below shows feature-related messages that you might see if you were downgrading from v5.2.0 to v5.1.0:

```
The following items need to be addressed before downloading the specified firmware:
Port mirror connections detected. Please use portmirror --delete to remove these
mirror connections.
AD feature is in use. Please clear it using the ad --clear command.
Port configuration with EX ports enabled along with trunking for port(s) 58, use the
portcfgexport, portcfgvexport, and/or portcfgtrunkport commands to disable the port
configuration. Verify that the blade is ENABLED. (error 3)
```

This example shows hardware-related messages for the same downgrade example:

```
switch:admin> firmwaredownload
Server Name or IP Address: 10.1.2.3
FTP User Name: JaneDoe
File Name: /pub/v5.1.0
FTP Password: xxxx
Verifying the input parameters ...
Checking system settings for firmwaredownload...

The following items need to be addressed before downloading the specified firmware:

AP BLADE type 31 is inserted. Please use slotshow to find out which slot it is in and
remove it.
SW BLADE type 36 is inserted. Please use slotshow to find out which slot it is in and
remove it.
Firmwaredownload command failed.

switch:admin>
```

Message *AP Blade type 24 is inserted. Please use slotshow to find out which slot it is in and remove it.*

Probable Cause and Recommended Action

The firmware download operation was attempting to download Fabric OS v5.0.0 with one or more SilkWorm FR4-18i port blades (blade ID 24) in the system. SilkWorm FR4-18i port blades are not supported on firmware v5.0.0 or lower, so the firmware download operation is aborted.

Use the **slotShow** command to display which slot the SilkWorm FR4-18i port blade is in, and physically remove the blade(s) from the chassis Retry the firmware download operation.

Message *Only platform options 1, 2, 5 are supported by version 5.1. Use chassisconfig to reset the option before downloading the firmware.*

Probable Cause and Recommended Action

The firmware download operation was attempting to upgrade a system to Fabric OS v5.1.0. The **chassisConfig** option was set to 3 or 4, which is not supported in v5.1.0, so the firmware download operation was aborted.

Execute the **chassisConfig** command with a supported option (1, 2, or 5 for SilkWorm 24000 and 1 or 5 for SilkWorm 48000), and then retry the firmware download operation. The supported options are described briefly below:

option 1	One 128-port switch with the following configuration: FC2-16 (blade ID 4), FC4-16 (blade ID 17) on slots 1–4 and 7–10; CP2 (blade ID 5), CP4 (blade ID 16) on slots 5–6
option 2	Two 64-port switches with the following configuration: FC2-16 (blade ID 4) on slots 1–4 and 7–10; CP2 (blade ID 5) on slots 5–6)
option 5	One 256-port switch with the following configuration: FC4-16 (blade ID 17), FC4-32 (blade ID 18) on slots 1–4 and 7–10; CP4 (blade ID 16) on slots 5–6

Message *Cannot download to 5.1 because Device Based routing policy is not supported by 5.1. Use aptPolicy to change the routing policy before proceeding.*

Probable Cause and Recommended Action

The firmware download operation was attempting to upgrade a system to Fabric OS v5.1.0 with device-based routing policy selected. Device-based routing policy is not supported in firmware v5.1.0, so the firmware download operation was aborted.

Disable the switch and change the routing policy selection to one of the following supported selections on firmware v5.1.0 using the **aptPolicy** command, and then retry the firmware download operation. The supported selections are described briefly below:

policy 1 Port-based routing policy

With this policy, the path chosen for an ingress frame is based on:
1. Ingress port on which the frame was received
2. Destination domain for the frame

The chosen path remains the same if Dynamic Load Sharing (DLS feature is not enabled. If DLS is enabled, then a different path might be chosen on a fabric event. Refer to the **dlsSet** command for the definition of a fabric event.

This policy may provide better ISL utilization when there is little or no oversubscription of the ISLs.

NOTE: Static routes are supported only with this policy.

policy 3 Exchange-based routing policy

With this policy, the path chosen for an ingress frame is based on:
1. Ingress port on which the frame was received
2. FC address of the SID for this frame
3. FC address of the DID for this frame
4. FC Originator Exchange ID (OXID) for this frame

This policy allows for optimal utilization of the available paths as I/O traffic between different (SID, DID, OXID) pairs can use different paths. All frames received on a ingress port with the same (SID, DID, OXID) parameters takes the same path unless there is a fabric event. Refer to the **dlsSet** command for the definition of a fabric event.

This policy does not support static routes. DLS always is enabled and the DLS setting cannot change with this policy.

Message *The command failed due to presence of long-distance ports in LS mode. Please remove these settings before proceeding.*

Probable Cause and Recommended Action

The firmware download operation was attempting to downgrade a system to Fabric OS v5.0.0 or lower with long-distance ports in LS mode. Long-distance ports in LS mode is not supported in firmware v5.0.0 or lower, so the firmware download operation failed.

Change the long distance port setting to a supported distance setting using the **portCfgLongDistance** command (the numerical value representing each distance level is shown in parentheses), and then retry the firmware download operation. The supported settings are described briefly below:

L0 (0) Specify L0 to configure the port to be a regular switch port. A total of 20 full-size frame buffers are reserved for data traffic, regardless of the port's operating speed; therefore, the maximum supported link distance is 10 km, 5 km, or 2.5 km for the port at speeds of 1 Gbit/sec, 2 Gbit/sec, or 4 Gbit/sec respectively.

L0.5 (4) Specify L0.5 (**portCfgShow** displays the two-letter code as LM) long distance, up to 25 km. A total of 12, 25, or 50 full-size frame buffers are reserved for data traffic for the port at speeds of 1 Gbit/sec, 2 Gbit/sec, or 4 Gbit/sec respectively.

L1 (1) Specify L1 long distance to support a long distance link up to 50 km. A total of 25, 50, or 100 full-size frame buffers are reserved for data traffic for the port at speeds of 1 Gbit/sec, 2 Gbit/sec, or 4 Gbit/sec respectively.

L2 (2) Specify L2 long distance to support a long distance link up to 100 km. A total of 50, 100, or 200 full-size frame buffers are reserved for data traffic for the port at speeds of 1 Gbit/sec, 2 Gbit/sec, or 4 Gbit/sec respectively. For previously released switches (Bloom1-based), the number of frames buffers is limited to 63.

LE (3) Specify LE mode is for E_Ports for distances beyond 5 Km and up to 10 Km. A total of 5, 10, or 20 full-size frame buffers are reserved for port speeds of 1 Gbit/sec, 2 Gbit/sec, or 4 Gbit/sec, respectively. LE does not require an Extended Fabrics license.

LD (5) Specify LD for automatic long-distance configuration. The buffer credits for the given E_Port are automatically configured, based on the actual link distance. Up to a total of 250 full-size frame buffers are reserved, depending upon the distance measured during E_Port initialization. If the desired distance is provided, it is used as the upper limit to the measured distance. For Bloom1-based systems, the number of frame buffers is limited to 63.

LS (6) Specify LS mode to configure a long-distance link with a fixed buffer allocation. Up to a total of 250 full-size frame buffers are reserved for data traffic, depending on the desired distance value provided with the **portCfgLongDistance** command. For Bloom1-based systems, the number of frame buffers is limited to 63.

Message *The command failed due to one or more ports having both long-distance and ISL R_RDY Modes enabled. Use portcfglongdistance and portcfgislmode to disable it before proceeding.*

Probable Cause and Recommended Action

The firmwareDownload operation was attempting to downgrade a system to Fabric OS v.0.0 or lower with both long-distance and ISL R_RDY modes enabled.

The **portCfgISLMode** and **portCfgLongDistance** modes cannot both be enabled at the same time; otherwise, fabric segmentation occurs, so the firmware download operation failed.

Disable either the long-distance on the port using the **portCfgLongDistance** command, or the ISL R_RDY mode using the **portCfgISLMode** command on the port or ports. Retry the firmware download operation.

Message *The command failed due to one or more ports having EX/VEX port configuration enabled. Use portcfgexport/portcfgvexport to disable it before proceeding.*

Probable Cause and Recommended Action

The firmware download operation was attempting to downgrade a system to Fabric OS v5.0.0 or lower with ports with EX or VEX configuration enabled. Ports with EX or VEX configuration enabled are not supported in firmware v5.0.0 or lower, so the firmware download operation failed.

Disable either the EX port configuration using the **portCfgExPort** command, or the VEX port configuration using the **portCfgVexPort** command. Retry the firmware download operation.

Message *Cannot downgrade to version 4.2 or lower. Please downgrade to 5.0 or 4.4 first and then download the desired version.*

Probable Cause and Recommended Action

The firmware download operation was attempting to downgrade a system from Fabric OS v5.1.0 directly to firmware v4.2.0 or lower. This firmware jump is not supported, so the firmware download operation aborted.

Downgrade the system to firmware v4.4.0 or v5.0.0 first, and then downgrade to the desired firmware version.

Message *The command failed due to network timeout.*

Probable Cause and Recommended Action

The firmware download operation was attempting to upgrade a system from Fabric OS v4.2.0 or lower directly to firmware v5.1.0. This firmware jump is not supported, so the firmware download operation aborted.

Note that the message is from the currently running switch firmware, which may be v4.2.0 or lower.

Upgrade the system to firmware v4.4.0 or v5.0.0 first, then upgrade to firmware v5.1.0.

Message *Cannot downgrade due to presence of port mirror connections. Use portmirror --delete to remove these mirror connections before proceeding.*

Probable Cause and Recommended Action

The firmware download operation was attempting to downgrade a system to Fabric OS v5.1.0 or lower with Port Mirroring enabled. Port Mirroring is not supported on firmware v5.1.0 or lower, so the firmware download operation failed.

Remove the mirror connections using the **portMirror - -delete** command. Retry the firmware download operation.

Message *AP Blade type 31 is inserted. Please use slotshow to find out which slot it is in and remove it.*

Probable Cause and Recommended Action

The firmware download operation was attempting to downgrade a system to Fabric OS v5.1.0 or lower with one or more SilkWorm FC4-16IP port blades (blade ID 31) in the system. SilkWorm FC4-16IP port blades are not supported on firmware v5.1.0 or lower, so the firmware download operation failed.

Use the **slotShow** command to display which slot the SilkWorm FC4-16IP port blade is in. Physically remove the blade(s) from the chassis, or use the micro-switch to turn the blade off. Retry the firmware download operation.

Message *SW Blade type 36 is inserted. Please use slotshow to find out which slot it is in and remove it.*

Probable Cause and Recommended Action

The firmware download operation was attempting to downgrade a system to Fabric OS v5.1.0 or lower with one or more SilkWorm FC4-48 port blades (blade ID 36) in the system. SilkWorm FC4-48 port blades are not supported on firmware v5.1.0 or lower, so the firmware download operation failed.

Use the **slotShow** command to display which slot the SilkWorm FC4-48 port blade is in. Physically remove the blade(s) from the chassis, or use the micro-switch to turn the blade off. Retry the firmware download operation.

Message *The command failed due to the presence of an Admin Domain. Use the ad command to remedy this before proceeding.*

Probable Cause and Recommended Action

The firmware download operation was attempting to downgrade a system to Fabric OS v5.1.0 or lower with Admin Domain (AD) enabled on the system. Admin Domains are not supported on firmware v5.1.0 or lower, so the firmware download operation failed.

Delete all of the Admin Domains using the one of the following **ad** commands:

ad --delete *ad_id* **ad --apply**	Deletes an Admin Domain. This command succeeds if the Admin Domain is in a deactivated or an activated state. Deleting AD1 through AD254 fails if the Admin Domain zone database is not empty. AD0 always exists; using this operation on AD0 only removes all the explicit members of AD0. AD0 zone database need not be empty before the delete operation. No existing user sessions to a deleted Admin Domain is terminated; however, subsequent CLI execution is disallowed. A message is displayed to indicate that the current Admin Domain is not active. **ad--delete** prompts for confirmation before deleting. The delete operation remains in the transaction buffer until you issue **ad --apply** or **ad --save**.
ad --clear *ad_id* **ad --apply**	Deletes all ADs and clears AD0's explicit members. This command fails if AD1 through AD254 zone database is not empty. **ad --clear** prompts for confirmation before deleting all ADs. The clear operation remains in the transaction buffer until you issue **ad --apply** or **ad --save**.

Retry the firmware download operation.

Message *The command failed due to the current zone size is not supported by the new firmware. Reduce the size of the configuration before proceeding.*

Probable Cause and Recommended Action

The firmware download operation was attempting to downgrade a system to Fabric OS v5.1.0 or lower and the current zone size is not supported by the firmware version to be downloaded, so the firmware download operation failed.

Reduce the zone database size to 256 KB. Verify that the zone size is below the 256 KB limit using the **cfgSize** command. Retry the firmware download operation.

Message *The command failed because Fast Write is enabled. Please use the portcfg fciptunnel command to disable it before proceeding.*

Probable Cause and Recommended Action

The firmware download operation was attempting to downgrade a system to Fabric OS v5.1.0 or lower and the fast write feature is enabled. Fast write is not supported on firmware v5.1.0 or lower, so the firmware download operation failed.

Disable the fast write feature using the **portCfg fcipTunnel** command. Retry the firmware download operation.

Message *The command failed because IPSec is enabled. Please use the portcfg fciptunnel command to disable it before proceeding.*

Probable Cause and Recommended Action

The firmwareDownload operation was attempting to downgrade a system to Fabric OS v5.1.0 or lower and the IPsec feature is enabled. The IPsec feature is not supported on firmware v5.1.0 or lower, so the firmwareDownload operation failed.

Disable IPSec using the **portCfg fcipTunnel** command. Retry the firmware download operation.

Message *The command failed because Tape Pipelining is enabled. Please use the portcfg fciptunnel command to disable it before proceeding.*

Probable Cause and Recommended Action

The firmware download operation was attempting to downgrade a system to Fabric OS v5.1.0 or lower and the tape pipelining feature is enabled. Tape pipelining is not supported on firmware v5.1.0 or lower, so the firmware download operation failed.

Disable tape pipelining using the **portCfg fcipTunnel** command. Retry the firmware download operation.

Message *The command failed due to the presence of trunking on an EX-Port. Please run switchCfgTrunk or portCfgTrunkPort CLI commands to disable the EX-Port trunking before proceeding.*

Probable Cause and Recommended Action

The firmware download operation was attempting to downgrade a system to Fabric OS v5.1.0 or lower and trunking is enabled on an EX_Port. EX_Port trunking is not supported on firmware v5.1.0 or lower, so the firmware download operation failed.

Disable the trunking on the EX_Port using the **portCfgTrunkPort** command, or disable trunking on all ports on the switch using the **switchCfgTrunk** command. Retry the firmware download operation.

Note

The messages might be displayed a number of times during the checking and summary phases of the firmware download operation. You need pay attention only to messages displayed *during the summary phase*.

Message *The command failed due to strict fabric wide policy is configured. Please use the fddcfg -fabwideset command to disable it before proceeding.*

Probable Cause and Recommended Action

The firmware download operation was attempting to downgrade a system to Fabric OS v5.1.0 or lower and a strict fabric-wide policy is enabled. The fabric-wide policy can be set to “strict” *only* when all switches in the fabric are running firmware v5.2.0 or higher, so the firmware download operation failed.

Either reset or disable the fabric-wide policy setting as follows:

- Reset the fabric-wide policy as “tolerant” using the **fddCfg --fabWideSet "policyname"** command. The “tolerant” setting allows switches with different firmware versions within the fabric, however all fabric-wide consistency mechanisms will be applied only to switches running firmware v5.2.0 or higher.
- Disable the strict fabric-wide policy using the **fddCfg --fabWideSet ""**command. The “absent” setting disables the fabric-wide consistency policy.

Retry the firmware download operation.

Message *The switch is currently configured with "radiuslocal" mode. Please use the aaaconfig command to remedy it before proceeding.*

Probable Cause and Recommended Action

The firmware download operation was attempting to downgrade a system to Fabric OS v5.1.0 or lower and radiuslocal mode is enabled. The radiuslocal mode is not supported on firmware v5.1.0 or lower, so the firmware download operation failed.

Disable the radiuslocal mode using the **aaaConfig -localOnly** command. Retry the firmwareDownload operation.

Message *The command failed due to DHCP is enabled. Please use the ipaddrset command to disable it before proceeding.*

Probable Cause and Recommended Action

The firmwareDownload operation was attempting to downgrade a system to Fabric OS v5.1.0 or lower and the DHCP feature is enabled. The DHCP feature is not supported on firmware v5.1.0 or lower, so the firmwareDownload operation failed.

Disable the DHCP feature using the **ipAddrSet** command. Retry the firmware download operation.

Message *Cannot downgrade directly to version 4.4 or lower. Please downgrade to 5.1 or 5.0 first and then download the desired version.*

Probable Cause and Recommended Action

The firmware download operation was attempting to downgrade a system directly to Fabric OS v4.4.0 or lower. This firmware jump is not supported, so the firmware download operation aborted.

Downgrade the system to firmware v5.1.0 or 5.0.0 first, and then download to the desired firmware version.

Message *The command failed due to network timeout.*

Probable Cause and Recommended Action

The firmware download operation was attempting to upgrade the system from Fabric OS v4.4.0 or lower directly to firmware v5.2.0. This firmware jump is not supported, so the firmware download operation aborted.

Upgrade the system to firmware v5.1.0 or 5.0.0 first, and then download firmware v5.2.0.

Message *DCC policy contains more than 256 numbers of ports. Please use the secpolicydelete and secpolicyactivate commands to remedy this.*

Probable Cause and Recommended Action

The firmware download operation was attempting to downgrade the system to Fabric OS v5.1.0 or lower and the number of ports in a DCC policy of the Active or Defined database is greater than 256. This DCC policy is not supported on firmware v5.1.0 or lower, so the firmwareDownload operation aborted.

Remove all DCC policies containing more than 256 ports using the **secPolicyDelete** and **secPolicyActivate** commands. Retry the firmware download operation.

Blade Troubleshooting Tips

Typically, issues detected during firmware download to AP blades do not require recovery actions on your part. However, if any of the following events occur, perform the suggested action to correct:

- The blade is faulty (issue **slotShow** to confirm).

  If the blade is faulty, enter the **slotPowerOff** and **slotPowerOn** commands for the blade. If the blade still appears to be faulty, remove it and re-insert it into the chassis.

- The blade is stuck in the “LOADING” state (issue **slotShow** to confirm).

  If the blade remains in the loading state for a significant period of time, the firmware download will time out. Remove the blade and re-insert it. When it boots up, autoleveling will be triggered and the firmware download will be attempted again.

If you experience frequent failovers between CPs that have different versions of firmware, then you may notice multiple blade firmware downloads and a longer startup time.

Synchronizing Firmware Versions on Partitions

It is possible for firmware to become out of sync on the partitions. This can occur when you issue the **firmwareDownload -s** command to test drive the firmware, but do not issue the **firmwareRestore** and **firmwareCommit** commands—which results in different levels of firmware on the two partitions.

If you have different versions of firmware on one of the partitions, for example, CP0 is running v5.0.1 on the primary and secondary partitions, and CP1 is running v5.0.1 on the primary partition and v4.4.0e on the secondary partition, then synchronize the partitions on CP1 as follows:

1. Start a telnet session on the CP with the out-of-sync partitions.
2. Enter the **firmwareCommit** command, which will copy the primary partition to the secondary partition.

If there is a discrepancy in the number of switches or attached devices in the fabric after a firmware download, it may be due to a fabric segmentation or a device that failed to log in. For assistance, refer to the following:

- Chapter 15, “Troubleshooting” for basic troubleshooting steps.
- Contact your switch service provider.

FTP Server Recommendations

The **firmwareDownload** command has a timeout of 10 minutes for v4.0 and v4.1 and 30 minutes for v4.2 and later. The following tips may be of help if you are having problems with the firmware download timing out before the process is complete:

- Verify that you can ping the FTP server from another host on the LAN, prior to running the **firmwareDownload** command.
- Disable any VPN, firewall or anti-virus applications on the FTP server.
- Remove any Ethernet hub between the FTP server and the switch.
- Optionally, set the switch's autonegotiating Ethernet port to match the speed and duplex setting of the IP environment.

  For example, enter the **ifModeSet** command to change the switch settings to match both ends of the Ethernet link. Enter **help ifModeSet** and **help ifModeShow** on the switch to view more details about these commands.
- Enable session logging on the FTP server so that the log can be viewed during the execution of the **firmwareDownload** command.
- Set the FTP server's timeout value to 60 minutes.
- Ensure that the decompress process created multiple *SWBDxx* folders (where *xx* is a number) in the main folder. If the files are unpacked without folder creation, then the **firmwareDownload** command will be unable to locate the *.plist* file.

Configuring SilkWorm Directors

Chapter 9

This chapter contains procedures that are specific to SilkWorm 24000 and 48000 directors.

Because directors contain interchangeable port blades (16/32/48-port blades in the SilkWorm 48000, as well as the FR4-18i, which is an 18-port blade and the FC4-16IP, which is a 16-port FC and iSCSI blade for the SilkWorm 48000), their procedures differ from those for SilkWorm 200E, 3250, 3850, 3900, 4100, 4900, and 7500 fixed-port switches. For example, fixed-port models identify ports by *domain, port* number, while director models identify ports by *slot/port* number.

SilkWorm 24000 and 48000 directors in their default configurations have only one domain (SilkWorm 48000 supports one domain only).

For detailed information about the SilkWorm 24000 and 48000 models, see the *SilkWorm 24000 Hardware Reference Manual* and the *SilkWorm 48000 Hardware Reference Manual*.

- "Identifying Ports" on page 9-1
- "Basic Blade Management" on page 9-7
- "Setting Chassis Configuration Options" on page 9-12
- "Setting the Blade Beacon Mode" on page 9-17

Identifying Ports

SilkWorm 24000 and 48000 models have slots and can have a variable number of ports within a given domain. Ports are identified by their combined slot number and port number.

There are 10 slots that contain port blades:

- Slot numbers 5 and 6 contain *control processor* blades (CPs).
- Slot numbers 1 through 4 and 7 through 10 contain *port* blades.

Note

CPs perform blade-to-blade switching over the director's passive backplane, contain communication ports for system management, and are used for low-level, chassis-wide tasks.

Port blades are used for host, storage, and interswitch connections.

On each port blade, there are 16 FC, 16 (8 FC and 8 iSCSI), 18 (16 FC ports and 2 GbE ports), 32, or 48 ports (counted from the bottom, 0 to 15, 0 to 31, or 0 to 23 and 24 to 47). A particular port must be represented by both slot number (1 through 4 and 7 through 10) and port number (0 through 15). For the

When you have port blades with different port counts in the same director (for example, 16-port blade and 32-port blades, or 16-port blades and 18-port blades with 16 FC ports and 2 GbE ports, or 16-port and 48-port blades), the area IDs no longer match the port numbers. Following are the port numbering schemes for the SilkWorm 48000:

- For the FC4-16 port blade, ports are numbered from 0 through 15 from bottom to top.
- For the FC4-32 port blade, ports are numbered from 0 through 15 from bottom to top on the left set of ports and 16 through 31 from bottom to top on the right set of ports.
- **SilkWorm 48000 only:** For the FR4-18i port blade, ports are numbered from 0 through 15 from bottom to top. There are also 8 GbE ports (numbered ge0 - ge7). So going from bottom to top, the 8 FC ports appear on the bottom, followed by the GbE ports at the top.
- **SilkWorm 48000 only:** For the FR4-16IP port blade, Fibre Channel ports are numbered from 0 through 7 from bottom to top. There are also 2 GbE ports (numbered ge0 - ge1). So going from bottom to top, the 2 physical GbE ports appear on the bottom of the blade followed by 16 physical FC ports.
- **SilkWorm 48000 only:** For the FR4-48 port blade, ports are numbered from 0 through 23 from bottom to top on the left set of ports and 24 through 47 from bottom to top on the right set of ports.

In the SilkWorm 48000, all the ports are part of a single logical switch. With Fabric OS v4.4.0 and later, you can configure the SilkWorm 24000 as two logical switches (*domains*).

The following sections tell how to identify ports on SilkWorm 24000 and 48000 models, and how to identify ports for zoning commands.

By Slot and Port Number

The port number is a number assigned to an external port to give it a unique identifier in a switch.

To select a specific port in the SilkWorm 24000, and 48000 models, you must identify both the slot number and the port number using the format *slot number*/*port number*. No spaces are allowed between the slot number, the slash (/), and the port number.

The following example shows how to enable port 4 on a blade in slot 2:

```
switch:admin> portenable 2/4
```

By Port Area ID

The relationship between the port number and area ID depends upon the PID format used in the fabric:

- When Core PID format is in effect, the area ID for port 0 is 0, for port 1, it is 1 and so forth.

  When using Core PID mode on the SilkWorm 24000 configured with two domains, the area IDs for both logical switches (domains) range from 0 to 63. This means that both logical switch 0 and logical switch 1 have a port that is referenced with area ID 0.

  For 32-port blades in the SilkWorm 48000 (using FC4-32), the numbering is contiguous up to port 15; from port 16, the numbering is still contiguous, but you must add 128 to each port number. For example, port 16 in slot 1 has a port number and area ID of 128; port number 15 has a port number and area ID of 15.

For 48-port blades in the SilkWorm 4800 (using FC4-48), the numbering is contiguous up to port 15; from port 16 the numbering is still contiguous, but you must add 128 to each port number. For example, port 48 in slot 1 has a port number and area ID of 176; port number 48 has a port number and area ID of 48.

- When Extended-edge PID format is in effect, the area ID is the port number plus 16 for ports 0 to 111. For port numbers higher than 111, the area ID wraps around so that port 112 has an area ID of 0, and so on. Each 64-port logical switch (domain) has area IDs ranging from 16 to 79. Port numbers higher than 128 are mapped directly to the core PID.

  For details about port area IDs in Extended-edge PID mode, see "Changing to Extended Edge PID Format" on page A-10.

- If you perform a port swap operation, the port number and area ID no longer match.

To determine the area ID of a particular port, enter the **switchShow** command. This command displays all ports on the current (logical) switch and their corresponding area IDs.

By Index

With the introduction of 48-port blades, the Index was introduced. Unique area IDs are possible up to 255 areas, but beyond that there needed to be some way to ensure uniqueness.

A number of fabric-wide databases supported by Fabric OS (including ZoneDB, Secure FOS DDC policies, the ACL DDC, and Admin Domain) allow a port to be designated by the use of a "D,P" (Domain, Port) notation. While the "P" component appears to be the port number, in up to 255 ports it was actually the *area* assigned to that port.

If the PID format is changed from Core to Extended-edge or vice-versa, the "P" value for ports 0-127 also changes. If two ports are changed using the portSwap command, their respective areas and "P" values are exchanged.

For ports that are numbered above 255, the "P" value is actually a logical index. The existing index-to-area_ID mapping for the first 256 ports in a switch is kept, and add values equal to the physical port number for the new FC4-48 ports.

The first 256 ports continue to have an index value equal to the area_ID assigned to the port. If a switch is using Core PID format, and no port swapping has been done, the port Index value for all ports is the same as the physical port numbers. Changing the PID format will change the index values as well as the area_ID values of the first 128 ports, and using portswap on a pair of ports will exchange those ports' area_ID and index values. Portswap is not supported for ports above 256.

Table 9-1 and Table 9-2 show the area ID and index mapping for core and extended-edge PID assignment. Note that up to 255 areas, the area_ID mapping to the index is one-to-one. Beyond this, the index is similar but not exact, and in some instances the area ID is shared among multiple ports.

These tables provide the area_ID/Index assignment for the maximum number of ports (used by the FC4-48 blade). If your blade does not have the maximum number of ports, use the lower sections of the table to determine the area_ID and index.

Table 9-1 Default Index/Area_ID Core PID Assignment with No Port Swap

Port on Blade	Slot 1Idx/ Area	Slot 2Idx/ Area	Slot 3Idx/ Area	Slot 4Idx/ Area	Slot 7Idx/ Area	Slot 8Idx/ Area	Slot 9Idx/ Area	Slot 10Idx/ Area
47	271/135	287/151	303/167	319/183	335/199	351/215	367/231	383/247
46	270/134	286/150	302/166	318/182	334/198	350/214	366/230	382/246
45	269/133	285/149	301/165	317/181	333/197	349/213	365/229	381/245
44	268/132	284/148	300/164	316/180	332/196	348/212	364/228	380/244
43	267/131	283/147	299/163	315/179	331/195	347/211	363/227	379/243
42	266/130	282/146	298/162	314/178	330/194	346/210	362/226	378/242
41	265/129	281/145	297/161	313/177	329/193	345/209	361/225	377/241
40	264/128	280/144	296/160	312/176	328/192	344/208	360/224	376/240
39	263/143	279/159	295/175	311/191	327/207	343/223	359/239	375/255
38	262/142	278/158	294/174	310/190	326/206	342/222	358/238	374/254
37	261/141	277/157	293/173	309/189	325/205	341/221	357/237	373/253
36	260/140	276/156	292/172	308/188	324/204	340/220	356/236	372/252
35	259/139	275/155	291/171	307/187	323/203	339/219	355/235	371/251
34	258/138	274/154	290/170	306/186	322/202	338/218	354/234	370/250
33	257/137	273/153	289/169	305/185	321/201	337/217	353/233	369/249
32	256/136	272/152	288/168	304/184	320/200	336/216	352/232	368/248
31	143/143	159/159	175/175	191/191	207/207	223/223	239/239	255/255
30	142/142	158/158	174/174	190/190	206/206	222/222	238/238	254/254
29	141/141	157/157	173/173	189/189	205/205	221/221	237/237	253/253
28	140/140	156/156	172/172	188/188	204/204	220/220	236/236	252/252
27	139/139	155/155	171/171	187/187	203/203	219/219	235/235	251/251
26	138/138	154/154	170/170	186/186	202/202	218/218	234/234	250/250
25	137/137	153/153	169/169	185/185	201/201	217/217	233/233	249/249
24	136/136	152/152	168/168	184/184	200/200	216/216	232/232	248/248
23	135/135	151/151	167/167	183/183	199/199	215/215	231/231	247/247
22	134/134	150/150	166/166	182/182	198/198	214/214	230/230	246/246
21	133/133	149/149	165/165	181/181	197/197	213/213	229/229	245/245
20	132/132	148/148	164/164	180/180	196/196	212/212	228/228	244/244
19	131/131	147/147	163/163	179/179	195/195	211/211	227/227	243/243
18	130/130	146/146	162/162	178/178	194/194	210/210	226/226	242/242

Table 9-1 Default Index/Area_ID Core PID Assignment with No Port Swap (Continued)

Port on Blade	Slot 1Idx/ Area	Slot 2Idx/ Area	Slot 3Idx/ Area	Slot 4Idx/ Area	Slot 7Idx/ Area	Slot 8Idx/ Area	Slot 9Idx/ Area	Slot 10Idx/ Area
17	129/129	145/145	161/161	177/177	193/193	209/209	225/225	241/241
16	128/128	144/144	160/160	176/176	192/192	208/208	224/224	240/240
15	15/15	31/31	47/47	63/63	79/79	95/95	111/111	127/127
14	14/14	30/30	46/46	62/62	78/78	94/94	110/110	126/126
13	13/13	29/29	45/45	61/61	77/77	93/93	109/109	125/125
12	12/12	28/28	44/44	60/60	76/76	92/92	108/108	124/124
11	11/11	27/27	43/43	59/59	75/75	91/91	107/107	123/123
10	10/10	26/26	42/42	58/58	74/74	90/90	106/106	122/122
9	9/9	25/25	41/41	57/57	73/73	89/89	105/105	121/121
8	8/8	24/24	40/40	56/56	72/72	88/88	104/104	120/120
7	7/7	23/23	39/39	55/55	71/71	87/87	103/103	119/119
6	6/6	22/22	38/38	54/54	70/70	86/86	102/102	118/118
5	5/5	21/21	37/37	53/53	69/69	85/85	101/101	117/117
4	4/4	20/20	36/36	52/52	68/68	84/84	100/100	116/116
3	3/3	19/19	35/35	51/51	67/67	83/83	99/99	115/115
2	2/2	18/18	34/34	50/50	66/66	82/82	98/98	114/114
1	1/1	17/17	33/33	49/49	65/65	81/81	97/97	113/113
0	0/0	16/16	32/32	48/48	64/64	80/80	96/96	112/112

Table 9-2 Default Index/Area Extended-edge PID Assignment with No Port Swap

Port on Blade	Slot 1Idx/ Area	Slot 2Idx/ Area	Slot 3Idx/ Area	Slot 4Idx/ Area	Slot 7Idx/ Area	Slot 8Idx/ Area	Slot 9Idx/ Area	Slot 10Idx/ Area
47	271/135	287/151	303/167	319/183	335/199	351/215	367/231	383/247
46	270/134	286/150	302/166	318/182	334/198	350/214	366/230	382/246
45	269/133	285/149	301/165	317/181	333/197	349/213	365/229	381/245
44	268/132	284/148	300/164	316/180	332/196	348/212	364/228	380/244
43	267/131	283/147	299/163	315/179	331/195	347/211	363/227	379/243
42	266/130	282/146	298/162	314/178	330/194	346/210	362/226	378/242
41	265/129	281/145	297/161	313/177	329/193	345/209	361/225	377/241
40	264/128	280/144	296/160	312/176	328/192	344/208	360/224	376/240
39	263/143	279/159	295/175	311/191	327/207	343/223	359/239	375/255
38	262/142	278/158	294/174	310/190	326/206	342/222	358/238	374/254
37	261/141	277/157	293/173	309/189	325/205	341/221	357/237	373/253

Table 9-2 Default Index/Area Extended-edge PID Assignment with No Port Swap (Continued)

Port on Blade	Slot 1Idx/ Area	Slot 2Idx/ Area	Slot 3Idx/ Area	Slot 4Idx/ Area	Slot 7Idx/ Area	Slot 8Idx/ Area	Slot 9Idx/ Area	Slot 10Idx/ Area
36	260/140	276/156	292/172	308/188	324/204	340/220	356/236	372/252
35	259/139	275/155	291/171	307/187	323/203	339/219	355/235	371/251
34	258/138	274/154	290/170	306/186	322/202	338/218	354/234	370/250
33	257/137	273/153	289/169	305/185	321/201	337/217	353/233	369/249
32	256/136	272/152	288/168	304/184	320/200	336/216	352/232	368/248
31	143/143	159/159	175/175	191/191	207/207	223/223	239/239	255/255
30	142/142	158/158	174/174	190/190	206/206	222/222	238/238	254/254
29	141/141	157/157	173/173	189/189	205/205	221/221	237/237	253/253
28	140/140	156/156	172/172	188/188	204/204	220/220	236/236	252/252
27	139/139	155/155	171/171	187/187	203/203	219/219	235/235	251/251
26	138/138	154/154	170/170	186/186	202/202	218/218	234/234	250/250
25	137/137	153/153	169/169	185/185	201/201	217/217	233/233	249/249
24	136/136	152/152	168/168	184/184	200/200	216/216	232/232	248/248
23	135/135	151/151	167/167	183/183	199/199	215/215	231/231	247/247
22	134/134	150/150	166/166	182/182	198/198	214/214	230/230	246/246
21	133/133	149/149	165/165	181/181	197/197	213/213	229/229	245/245
20	132/132	148/148	164/164	180/180	196/196	212/212	228/228	244/244
19	131/131	147/147	163/163	179/179	195/195	211/211	227/227	243/243
18	130/130	146/146	162/162	178/178	194/194	210/210	226/226	242/242
17	129/129	145/145	161/161	177/177	193/193	209/209	225/225	241/241
16	128/128	144/144	160/160	176/176	192/192	208/208	224/224	240/240
15	31/31	47/47	63/63	79/79	95/95	111/111	127/127	15/15
14	30/30	46/46	62/62	78/78	94/94	110/110	126/126	14/14
13	29/29	45/45	61/61	77/77	93/93	109/109	125/125	13/13
12	28/28	44/44	60/60	76/76	92/92	108/108	124/124	12/12
11	27/27	43/43	59/59	75/75	91/91	107/107	123/123	11/11
10	26/26	42/42	58/58	74/74	90/90	106/106	122/122	10/10
9	25/25	41/41	57/57	73/73	89/89	105/105	121/121	9/9
8	24/24	40/40	56/56	72/72	88/88	104/104	120/120	8/8
7	23/23	39/39	55/55	71/71	87/87	103/103	119/119	7/7
6	22/22	38/38	54/54	70/70	86/86	102/102	118/118	6/6
5	21/21	37/37	53/53	69/69	85/85	101/101	117/117	5/5

Table 9-2 Default Index/Area Extended-edge PID Assignment with No Port Swap (Continued)

Port on Blade	Slot 1Idx/ Area	Slot 2Idx/ Area	Slot 3Idx/ Area	Slot 4Idx/ Area	Slot 7Idx/ Area	Slot 8Idx/ Area	Slot 9Idx/ Area	Slot 10Idx/ Area
4	20/20	36/36	52/52	68/68	84/84	100/100	116/116	4/4
3	19/19	35/35	51/51	67/67	83/83	99/99	115/115	3/3
2	18/18	34/34	50/50	66/66	82/82	98/98	114/114	2/2
1	17/17	33/33	49/49	65/65	81/81	97/97	113/113	1/1
0	16/16	32/32	48/48	64/64	80/80	96/96	112/112	0/0

Basic Blade Management

The following sections provide procedures for powering a port blade off and on and for disabling and enabling a port blade.

Powering Port Blades Off and On

Port blades are powered on by default.

To power off a port blade

1. Connect to the switch and log in as admin.
2. Enter the **slotPowerOff** command with the slot number of the port blade you want to power off.

   The slot must exist in the logical switch where you are logged in.

```
switch:admin> slotpoweroff 3
Slot 3 is being powered off
switch:admin>
```

To provide power to a port blade

1. Connect to the switch and log in as admin.
2. Enter the **slotPowerOn** command with the slot number of the port blade you want to power on.

   The slot must exist in the logical switch where you are logged in.

```
switch:admin> slotpoweron 3
Powering on slot 3
switch:admin>
```

Disabling and Enabling Port Blades

Port blades are enabled by default. In some cases, you will need to disable a port blade to perform diagnostics. When diagnostics are executed manually (from the Fabric OS command line), many commands require the port blade to be disabled. This ensures that diagnostic activity does not interfere with normal fabric traffic.

To disable a port blade

1. Connect to the switch and log in as admin.
2. Enter the **bladeDisable** command with the slot number of the port blade you want to disable:

```
switch:admin> bladedisable 3
Slot 3 is being disabled
```

To enable a port blade

1. Connect to the switch and log in as admin.
2. Enter the **bladeEnable** command with the slot number of the port blade you want to enable:

```
switch:admin> bladeenable 3
Slot 3 is being enabled
```

SilkWorm 7500 Exceptions

The first time the SilkWorm 7500 is powered on ports are persistently disabled. Ports will remain disabled until they are configured otherwise.

FR4-18i Blade Exceptions

You may wish to persistently disable FR4-18i blade ports that are not configured so they cannot join the fabric when the following scenarios apply:

- You have inserted the FR4-18i blade into a slot that was previously empty or contained an FC4-48, FC4-32, FC4-16, or FC4-16IP.
- You have turned on the power of the chassis and the FR4-18i blade in that slot was not active prior to the power-on.

If the FR4-18i blade is operational and the switch is rebooted, then after the successful boot up of the system the blade will continue operations using the previous configurations.

If a previously configured FR4-18i blade is removed and another or the same FR4-18i blade is inserted into the same slot, then the ports will use the previous configuration and come up enabled.

If a previously-configured FR4-18i blade is removed and an FC4-48, FC4-32, or FC4-16 blade is plugged in, then–other than the port's EX_Port configuration–all the remaining port configurations previously applied to the FR4-18i FC_Ports can be used. So the EX_Port configuration on those ports will be disabled before the FC4-48, FC4-32, or FC4-16 blade becomes operational. When a blade is present in the slot, then any requested port configuration is validated against the blade's capabilities before accepting the request.

To summarize:

- When an FC4-16 or FC4-32 blade is replaced by an FR4-18i blade, the FC configuration of the previously configured FC_Ports will continue to be used, and all FC_Ports on the FR4-18i blade will be persistently disabled.
- When an FR4-18i blade is replaced by an FC4-16 or FC4-32 blade, then the EX_Port configuration will be removed from any ports that were configured as EX_Ports (equivalent to disabling the EX_Port configuration using the **portCfgEXPort** command). All remaining port configurations will be retained.

FC4-48 Blade Exceptions

The FC4-48 blade is compatible only with the FC4 control processor blade, FC4-16/32 port blade, FR4-18i, and FC4-16IP iSCSI blade.

Because the area IDs are shared with different port IDs, the FC4-48 blade only supports F, G and E_Ports.

Port swapping on an FC4-48 is supported only on ports 0–15. For the FC4-32 blade port swapping is supported on all 32 ports. This means that if you replace an FC4-32 where a port has been swapped on ports 16–31 with an FC4-48, the FC4-48 faults. To correct this, reinsert the FC4-32 and issue **portSwap** to restore the original area IDs to ports 16–31.

Conserving Power

To conserve power and ensure that more critical components are the least affected by a power fluctuation, you can power off components in a specified order, using the **powerOffListSet** command.

The available power is compared to the power demand to determine if there is enough power to operate. If there is less power available than the demand, the power-off list is processed until there is enough power for operation. By default, the processing proceeds from slot 1 to the last slot in the chassis. As power becomes available, slots are powered up in the reverse order.

Note

Some FRUs in the chassis may use significant power, yet cannot be powered off through software. For example, a missing blower FRU may change the power computation enough to affect how many slots can be powered up.

The **powerOffListShow** command displays the power off order.

Blade Terminology and Compatibility

Before configuring a chassis, familiarize yourself with the director CP blade and port blade nomenclature, as well as the port blade compatibilities. Often in procedures, only the abbreviated names for CP and port blades are used (for example, the FC4-16 blade). Table 9-3 includes CP and port blade abbreviations and descriptions.

Table 9-3 SilkWorm Director Terminology and Abbreviations

Term	Abbreviation	Blade ID (slotshow)	Definition
SilkWorm 24000 control processor blade	CP2	5	The second generation CP blade provided with the SilkWorm 24000. This CP supports 1, 2, and 4 Gbit/sec port speeds. It supports both the dual domain and a single domain configuration within the chassis.
SilkWorm 48000 control processor blade	CP4	16	The third generation CP blade provided with the SilkWorm 48000. This CP supports 1, 2, and 4 Gbit/sec port speeds, as well as 16 and 32-port blades.
16-port 2-Gbit/sec port blade	FC-16	2	The first generation SilkWorm director 16-port blade supporting 2 Gbit/sec port speeds. This port blade is only compatible with the SilkWorm 24000 CP blades.
16-port 2-Gbit/sec port blade	FC2-16	4	The second generation SilkWorm director 16 port blade supporting 1 and 2 Gbit/sec port speeds. This port blade is only compatible with the SilkWorm 24000 or SilkWorm 48000 CP blades.
16-port 4-Gbit/sec port blade	FC4-16	17	The third generation SilkWorm director 16 port blade supporting 1, 2, and 4 Gbit/sec port speeds. This port blade is only compatible with the SilkWorm 24000 or SilkWorm 48000 CP blades. FC4-16 blades do not support Fibre Channel, Arbitrated Loop private devices.
8-port 1GbE iSCSI ports and 8-port 4 Gbit/sec FC ports	FC4-16IP	31	
16-port 4-Gbit/sec port blade with 2-port, 1 GbE (FCIP) capabilities	FR4-18i	24	A16-port Fibre Channel routing and FCIP blade that also has 2 GbE ports and is compatible only with the SilkWorm 48000 (using chassis configuration option 5).
32-port 4-Gbit/sec port blade	FC4-32	18	A 32 port SilkWorm director port blade supporting 1, 2, and 4 Gbit/sec port speeds. This port blade is only compatible with the SilkWorm 48000 CP blades. FC4-32 blades do not support private devices.
48-port 4-Gbit/sec port blades	FC4-48	36	A 48 port SilkWorm director port blade supporting 1, 2, and 4 Gbit/sec port speeds in chassis mode 5 with port and exchange-based routing. This port blade is only compatible with the SilkWorm 48000 CP blades.

CP Blades

CP blades determine the director type:

- If CP2 blades are installed, the director is a SilkWorm 24000.
- If CP4 blades are installed, the director is a SilkWorm 48000.

Mixed CP blades are not supported on a single chassis, except during specific upgrade procedures detailed in the *SilkWorm 48000 Hardware Reference Manual*. Brocade recommends that each SilkWorm director have only one type of CP blade installed and that each CP (primary and secondary partition) maintains the same firmware version.

Port Blade Compatibility

Table 9-4 identifies which blades are supported for each SilkWorm director.

Table 9-4 Blades Supported by Each SilkWorm Director

Director	Port Blades						
	FC-16	FC2-16	FC4-16	FC4-16IP	FC4-32	FC4-48	FR4-18i
SilkWorm 24000 (CP2)	Supported	Supported	Supported		NA	NA	NA
SilkWorm 48000 (CP4)	NA	Supported[a]	Supported	Supported	Supported	Supported	Supported

a. Can coexist only with FC4-16 blades.

Setting Chassis Configuration Options

The **chassisConfig** command allows you to set the chassis configuration for the SilkWorm 24000, which supports both single domain and dual domain operation.

The SilkWorm 48000 allows you to use chassis configuration options 1 and 5. Configuration option 1 is compatible with FC4-16 blades; configuration option 5 is compatible with FC4-16IP, FC4-32, FR4-18i, and FC4-48 blades.

Table 9-5 lists the supported configurations options for Fabric OS v5.2.0.

Table 9-5 Supported Configuration Options

Option	Number of Domains	Maximum Number of Ports Per Switch	Supported Port Blades	Supported CP Blades	Notes
1	1	128	FC2-16, FC4-16	CP2 or CP4	Option 1 is the default configuration for SilkWorm 24000.
2	2	64/64	FC2-16	CP2	
5	1	384	FC4-16, FC4-16IP, FC4-32, FR4-18i, FR4-48[a]	CP4	Option 5 is the default configuration option for SilkWorm 48000.

a. L_Ports are not supported on the FC4-48 blade.

Table 9-6 lists configuration options and resulting slot configurations.

Table 9-6

Option	Result
1	One 128-port switch (Blade IDs 4, 17 on slots 1–4, 7–10. Blade ID 5 and 16 on slots 5, 6)
2	Two 64-port switches (Blade ID 4 on slots 1–4, 7–10. Blade ID 5 on slots 5, 6)
5	One 384-port switch (Blade IDs 4, 17, 18, 31, and 36 on slots 1–4, 7–10. Blade ID 16 on slots 5, 6)

See Table 9-3 on page 9-10 for details about the different blades, including their corresponding IDs.

The following sections contain procedures for obtaining chassis information and for configuring director domains using the **chassisConfig** command.

Obtaining Slot Information

For a SilkWorm 24000 configured as two logical switches, the chassis-wide commands display or control both logical switches. In the default configuration, SilkWorm 24000 directors are configured as one logical switch, so the chassis-wide commands display and control the single logical switch.

To display the status of all slots in the chassis

1. Connect to the switch and log in as user or admin.
2. Enter the **slotShow** command to display the current status of each slot in the system. The format of the display includes a header and four fields for each slot. The fields and their possible values are:

Slot	Displays the physical slot number.
Blade Type	Displays the blade type: SW BLADE: The blade is a switch. CP BLADE: The blade is a control processor. AP BLADE: The blade is the FR4-18i blade. UNKNOWN: The blade is not present or its type is not recognized.
ID	Displays the hardware ID of the blade type. See Table 9-3 on page 9-10 for a list of blades and their corresponding IDs.
Status	Displays the status of the blade: VACANT: The slot is empty. INSERTED, NOT POWERED ON: The blade is present in the slot but is turned off. POWERING UP: The blade is present and powering on. LOADING: The blade is present, powered on, and loading initial configuration. DIAG RUNNING POST1: The blade is present, powered on, and running the post-initialization power-on self test. DIAG RUNNING POST2: The blade is present, powered on, and running the POST (power-on self test). INITIALIZING: The blade is present, powered on, and initializing hardware components. ENABLED: The blade is on and enabled. ENABLED (User Ports Disabled): The blade is on, but external ports have been disabled with the **bladeDisable** command. DISABLED: The blade is powered on but disabled. FAULTY: The blade is faulty because an error was detected. The reason code numbers displayed are for debugging purposes. UNKNOWN: The blade is inserted but its state cannot be determined.

Configuring a New SilkWorm 24000 with Two Domains

By default, the SilkWorm 24000 director is configured as one 128-port switch (one domain). The procedure assumes that the new director:

- Has been installed and connected to power, but is not yet attached to the fabric.
- Has been given an IP address, but is otherwise running factory defaults.

  If this is not the case, back up the current configuration before starting, so that you can restore it later if necessary.
- Is running Fabric OS v4.4.0 or later.
- Is running in configuration option one (one switch, FC2-16 cards installed).

Use the following procedure to add a factory new SilkWorm 24000 to a fabric and configure it as two 64-port switches (two domains).

To add a new SilkWorm 24000 to a fabric and configure it

1. Connect to the switch and log in as admin.
2. Enter the **chassisConfig** command without options to verify that the switch is configured with one domain. For example:

```
switch:admin> chassisconfig
Current Option: 1
```

3. Enter the **chassisConfig** command to configure two domains. Use the **-f** option to suppress prompting for uploading the configuration. This command reboots the system.

```
switch:admin> chassisconfig -f 2
Current Option changed to 2
Restoring switch 0 configuration to factory defaults...
All account passwords have been successfully set to factory default.
Restoring switch 1 configuration to factory defaults...
All account passwords have been successfully set to factory default.
```

4. After the system reboots, log in again to the first logical switch (sw0) as admin.
5. Use the **configure** command to configure the sw0 to match your fabric specifications.

   If the director is to be merged into an existing fabric, do not configure zoning parameters; these will be propagated automatically when you merge the director into the fabric.
6. Log in to the second logical switch (sw1) as admin.
7. Use the **configure** command to configure the sw1 to match your fabric specifications.

   If the director is to be merged into an existing fabric, do not configure zoning parameters; these will be propagated automatically when you merge the director into the fabric.
8. If the fabric is in secure mode, perform the following steps; otherwise, proceed to step 9.

   (See the *Secure Fabric OS Administrator's Guide* for specific instructions.)

a. Optionally, to configure sw0 and sw1 in one operation, connect them with an ISL link to form a temporary fabric.

b. If you want sw0 and sw1 to be fabric configuration servers, update the overall fabric's FCS policy to include them. If not, skip this step.

c. On sw0, enable security mode and use the **secModeEnable** command to create an FCS list that matches your overall fabric's FCS policy.

d. Reset the version stamp on sw0.

e. If you connected sw0 and sw1 in step a and you do not want them connected, disconnect the ISL link between them. If you did not connect them, repeat step b through step d on sw1.

9. Optionally, connect the new two-domain SilkWorm 24000 director to the fabric.
10. Enter the **fabricShow** command to verify that sw0 and sw1 have been merged with the fabric.
11. Enter the **cfgShow** command to verify that zoning parameters were propagated.

Converting an Installed SilkWorm 24000 to Support Two Domains

Fabric OS versions earlier than v4.4.0 supported only one domain for SilkWorm 24000 directors (one 128-port logical switch). When you upgrade a SilkWorm 24000 director to Fabric OS v4.4.0 or later, you can use the **chassisConfig** command to specify two domains for the director (two 64-port logical switches, sw0 and sw1). This conversion is for SilkWorm 24000 directors using configuration option one (one switch, FC2-16 cards installed).

Note

This procedure restores most configuration parameters to factory defaults. After performing this procedure, you must check the new configuration and reconfigure those parameters that you customized in the old configuration.

During this procedure, power is reset and the CP blades are rebooted, so traffic on the fabric is disrupted. If the fabric is in secure mode, enabling security on the new domains is a complicated task. You should avoid converting existing SilkWorm 24000s that are acting as the core in a core-edge network.

To convert a SilkWorm 24000 to support two domains

1. Connect to the switch and log in as admin.
2. If the director is already in a fabric, minimize disruption by removing the director from the fabric using one of the following methods:
   - Physically disconnect the director.
   - Use the **portCfgPersistentDisable** command on all connected remote switches to persistently disable their ports that are connected to the director, or remove ISLs that connect the SilkWorm 24000 to the current fabric.
3. Enter the **chassisConfig** command to change the configuration from the default (one domain) to two domains. This command reboots the system.

   ```
   switch:admin> chassisconfig 2
   ```

   During the conversion, you are prompted to save the configuration of sw0. Follow the prompts to save the configuration file.
4. Enter the **ipAddrSet** command to set and confirm the IP address of sw1 (sw1 takes on a default that must be corrected). The IP address of sw0 will already be set.
5. After the system reboots, log in again as admin to each logical switch and enter the **switchName** command to assign a name to the new switch.
6. Using the configuration file saved in step 3 as a guide, manually reconfigure sw0 and sw1.

   Do not configure zoning parameters; these are propagated automatically when you merge the director into the fabric.
7. If the fabric is in secure mode, perform the following steps; otherwise, proceed to step 8.
   a. Optionally, to configure sw0 and sw1 in one operation, connect them with an ISL link to form a temporary fabric.
   b. If you want sw0 and sw1 to be fabric configuration servers, update the overall fabric's FCS policy to include them. If not, skip this step.
   c. On sw0, enable security mode and use the **secModeEnable** command to create an FCS list that matches your overall fabric's FCS policy.
   d. Reset the version stamp on sw0.
   e. If you connected sw0 and sw1 in step a and you do not want them connected, disconnect the ISL link between them. If you did not connect them, repeat step b through step d on sw1.
8. If you physically disconnected the switch in step 2, reconnect it to the fabric.

   If you used the **portCfgPersistentDisable** command in step 2, use the **portCfgPersistentEnable** command to persistently enable all ports that connect the switch to other switches in the fabric.
9. Enter the **fabricShow** command to verify that sw0 and sw1 have been merged with the fabric.
10. Enter the **configShow** command to verify that zoning parameters were propagated.

Setting the Blade Beacon Mode

When beaconing mode is enabled, the port LEDs will flash amber in a running pattern from port 0 through port 15 and back again. The pattern continues until the user turns it off. This can be used to locate a particular blade.

To set the blade beacon mode on

1. Connect to the switch and log in as admin.
2. Enter the **bladeBeacon** command:

```
switch:admin> bladebeacon slotnumber, mode
```

The *slotnumber* is the blade on which you want to enable beacon mode; this slot number must exist on the logical switch. The value 1 turns beaconing mode on, and 0 turns beaconing mode off.

```
switch:admin> bladebeacon 3, 1
```

Routing Traffic

Chapter 10

This chapter describes SilkWorm switch routing features and procedures.

- "About Data Routing and Routing Policies" on page 10-1
- "Specifying the Routing Policy" on page 10-2
- "Assigning a Static Route" on page 10-3
- "Specifying Frame Order Delivery" on page 10-3
- "Using Dynamic Load Sharing" on page 10-4
- "Viewing Routing Path Information" on page 10-5
- "Viewing Routing Information Along a Path" on page 10-7

About Data Routing and Routing Policies

Data moves through a fabric from switch to switch and storage to server along one or more paths that make up a *route*. Routing policies determine the correct path for each frame of data.

Whatever routing policy a switch is using applies to the VE_Ports as well. See Chapter 11, "Using the FC-FC Routing Service" for details about VE_Ports.

Caution
For most configurations, the default routing policy is optimal, and provides the best performance. You should only change the policy if there is a performance issue that is of concern, or a particular fabric configuration requires it.

The following routing policies are available to tune routing performance:

- Exchange-based routing

  The choice of routing path is based on the SID, DID, and Fibre Channel originator exchange ID (OXID), optimizing path utilization for the best performance. Thus, every exchange can take a different path through the fabric. Exchange-based routing requires the use of the Dynamic Load Sharing (DLS) feature; when these policies are in effect, you cannot disable the DLS feature.

- Port-based routing

  The choice of routing path is based only on the incoming port and the destination domain. To optimize port-based routing, DSL can be enabled to balance the load across the available output ports within a domain.

  Using port-based routing, you can assign a "static route", in which the path chosen for traffic never changes. In contrast, exchange-based routing policies always employ "dynamic path selection." Port-based routing is supported by all SilkWorm models.

Specifying the Routing Policy

The following routing policies are supported:

- Port-based path selection

  Default on SilkWorm 3250, 3850, 3900, 24000, and 48000 (using configuration option 1). These switches support the port-based policy only; you cannot change the routing policy for these switches. SilkWorm 200E, 4100, 4900, 7500, and 48000 (using configuration option 5) switches can also use port-based routing.

- Exchange-based path selection

  Default on SilkWorm 200E, 4100, 4900, 7500, and 48000 (using configuration option 5).

See Chapter 9, “Configuring SilkWorm Directors” for more details about SilkWorm 48000 configuration options.

You can use the **aptPolicy** command to display and specify a different routing policy. The **aptPolicy** command detects the switch’s configuration options and provides the appropriate policies for users to select from. For example, if you attempt to set the policy for the SilkWorm 48000 using configuration option 1, an error message is returned because you cannot change the routing policy. See the *Fabric OS Command Reference* for more details on the **aptPolicy** command.

You must disable the switch before changing the routing policy, and re-enable it afterward.

In the following example, the routing policy for a SilkWorm 7500 is changed from exchange-based to port-based:

```
switch:admin> aptpolicy
 Current Policy: 3

 3: Default Policy
 1: Port Based Routing Policy
 3: Exchange Based Routing Policy

switch:admin> switchdisable
switch:admin> aptpolicy 1
Policy updated successfully.
switch:admin> switchenable
switch:admin> aptpolicy
Current Policy: 1
```

Assigning a Static Route

A static route can be assigned only when the active routing policy is port-based and running on a SilkWorm 200E, 3250, 3850, 3900, 4100, 24000, and 48000 (using configuration option 1). When exchange-based routing is active, you cannot assign static routes. Even if the active routing policy is port-based, you cannot assign static routing to the SilkWorm 4900, 7500, or 48000 (using configuration option 5).

To assign a static route, use the **uRouteConfig** command. To remove a static route, use the **uRouteRemove** command

Note

SilkWorm 3900, 24000, and 48000 (using configuration option 1):

When you enter the **uRouteConfig** command, two similar warning messages might display if a platform conflict occurs. The first message displays when the static routing feature detects the conflict. The second message displays when the Dynamic Load Sharing feature detects the condition as it tries to rebalance the route.

A platform conflict occurs if a static route was configured with a destination port that is currently down. The static route is ignored in this case, in favor of a normal dynamic route. When the configured destination port comes back up, the system attempts to reestablish the static route, potentially causing a conflict.

Specifying Frame Order Delivery

The order of delivery of frames is maintained within a switch and determined by the routing policy in effect. Following are the frame delivery behaviors for each routing policy.

- Port-based routing

  All frames received on an ingress port destined for a destination domain are guaranteed to exit the switch in the same order in which they were received.

- Exchange-based routing

  All frames received on an ingress port for a given exchange are guaranteed to exit the switch in the same order in which they were received. Because different paths are chosen for different exchanges, this policy does not maintain the order of frames across exchanges.

If even one switch in the fabric delivers out-of-order exchanges, then exchanges are delivered to the target out-of-order, regardless of the policy configured on other switches in the fabric.

Note

Certain devices do not tolerate out-of-order exchanges; in such cases, use the port-based routing policy.

In a stable fabric, frames are always delivered in order, even when the traffic between switches is shared among multiple paths. However, when topology changes occur in the fabric (for example, if a link goes down), traffic is rerouted around the failure, and some frames could be delivered out of order. Most destination devices tolerate out-of-order delivery, but some do not.

By default, out-of-order frame-based delivery is allowed to minimize the number of frames dropped. You should only force in-order frame delivery across topology changes if the fabric contains destination devices that cannot tolerate occasional out-of-order frame delivery.

To force in-order frame delivery across topology changes

1. Connect to the switch and log in as admin.
2. At the command line, enter the **iodSet** command.

Note

This command can cause a delay in the establishment of a new path when a topology change occurs; use it with care.

3. To confirm the In Order Delivery has been set, issue the **iodShow** command.

To restore out-of-order frame delivery across topology changes

1. Connect to the switch and log in as admin.
2. Enter the **iodReset** command at the command line.

Using Dynamic Load Sharing

The exchange-based routing policy depends on the Fabric OS Dynamic Load Sharing feature (DLS) for dynamic routing path selection. When using the exchange-based routing policy, DLS is by default enabled and cannot be disabled. In other words, you cannot enable or disable DLS when the exchange-based routing policy is in effect.

When the port-based policy is in force, you can enable DLS to optimize routing. When DLS is enabled, it shares traffic among multiple equivalent paths between switches. DLS recomputes load sharing when a switch boots up or each time an E_Port goes offline and online, or a Fx_Port goes offline.

To check and set DLS

1. Connect to the switch and log in as admin.
2. Enter the **dlsShow** command to view the current DLS setting.

   One of the following messages appears:

   - "DLS is set" indicates that dynamic load sharing is turned on.
   - "DLS is not set" indicates that dynamic load sharing is turned off.
   - "DLS cannot be changed with current routing policy" indicates that you are using the exchange-based routing policy and you cannot enable or disable DLS.

3. Enter the **dlsSet** command to enable DLS or enter the **dlsReset** command to disable it.

```
switch:admin> dlsshow
DLS is not set
switch:admin> dlsset
switch:admin> dlsshow
DLS is set
switch:admin> dlsreset
switch:admin> dlsshow
DLS is not set
```

Viewing Routing Path Information

The **topologyShow** and **uRouteShow** commands provide information about the routing path.

1. Connect to the switch and log in as admin.
2. Enter the **topologyShow** command to display the fabric topology, as it appears to the local switch:

```
switch:admin> topologyshow

4 domains in the fabric; Local Domain ID: 2

Domain:         1
Metric:         10500
Name:           fcr_xd_1_1
Path Count:     1

        Hops:                   2
        Out Port:               39
        In Ports:               35 56
        Total Bandwidth:        4.000 Gbps
        Bandwidth Demand:       300 %
        Flags:                  D
switch:admin>
```

The following entries appear:

Local Domain ID	Domain number of the local switch
Domain	Domain number of the destination switch
Metric	Cost of reaching the destination domain
Name	The name of the destination switch
Path Count	The number of currently active paths to the destination domain
Hops	The maximum number of switch-to-switch links (ISLs) transversed to reach the destination domain
Out Port	The Port to which the incoming frame will be forwarded in order to reach the destination domain

In Ports	Input ports that use the corresponding Out Port to reach the destination domain
Total Bandwidth	The maximum bandwidth of the out port
Bandwidth Demand	The maximum bandwidth demand of the in ports
Flags	Always “D”, indicating a dynamic path

3. Use the **uRouteShow** command to display unicast routing information.

   SilkWorm 200E, 3250, 3850, 3900, 4100, 4900 and 7500: Use the following syntax:

   urouteshow [*portnumber*][, *domainnumber*]

   SilkWorm 24000 and 48000: Use the following syntax:

   urouteshow [*slot*/][*portnumber*][, *domainnumber*]

   The following entries appear:

Local Domain	Domain number of the local switch
In Ports	Port from which a frame is received
Domain	Destination domain of the incoming frame
Out Port	The port to which the incoming frame will be forwarded in order to reach the destination domain
Name	The name of the destination switch
Path Count	The number of currently active paths to the destination domain
Hops	The maximum number of switch-to-switch links (ISLs) transversed to reach the destination domain
Flags	Indicates if the route is dynamic (D) or static (S). A static route is assigned using the command **uRouteConfig**.
Next (Dom, Port)	Domain number and port number of the next hop

   The following example displays the routing information of all the active ports:

```
switch:admin> urouteshow

Local Domain ID: 1

In Port    Domain    Out Port    Metric    Hops    Flags    Next (Dom, Port)
-------------------------------------------------------------------------
  14          4          9         500       1       D            4,24

switch:admin>
```

   The next example displays the routing information for port 11 on slot 1:

```
switch:admin> urouteshow 1/11
Local Domain ID: 3
In Port    Domain    Out Port    Metric    Hops    Flags    Next (Dom, Port)
-------------------------------------------------------------------------
 11         2         0         1500      2       D            4,0
            4        16          500      1       D            4,0
```

This example displays the routing information of port 11 to domain 4 only:

```
switch:admin> urouteshow 1/11, 4
Local Domain ID: 3
In Port    Domain    Out Port    Metric    Hops    Flags    Next (Dom, Port)
------------------------------------------------------------------------------
 11         4        16          500      1       D            4,0
```

Viewing Routing Information Along a Path

You can display detailed routing information from a source port (or area) on the local switch to a destination port (or area) on another switch. This routing information describes the full path that a data stream travels between these ports, including all intermediate switches.

1. Connect to the switch and log in as admin.
2. Enter the **pathInfo** command. In interactive mode, you can specify the following parameters for display:

Max hops	The maximum number of hops that the pathinfo frame is allowed to traverse
Domain	The destination domain ID
Source Port	The port number (or area number for SilkWorm 24000 or 48000 directors) on which the switch receives frames
Destination Port	The output port that the frames use to reach the next hop on this path. For the last hop, the destination port.
Basic stats	Basic statistics on every link
Extended stats	Detailed statistics on every link
Trace reverse path	Traverses from the destination switch back to the source switches
Source route	Forces the frame to follow a specified path to reach the destination
Timeout	The maximum time to wait for a response from **pathInfo**, in seconds

Paths always originate on the local switch. The path destination can be specified by domain or port. By default, the path will be the path taken by traffic from the source to destination port, but you can also specify all or portions of a path.

See the *Fabric OS Command Reference* for details on the **pathInfo** command.

The following example is from a SilkWorm 3900 switch (other models provide similar information):

```
switch:admin> pathinfo

    Max hops: (1..127) [25]
    Domain: (1..239) [-1] 1
    Source port: (0..255) [-1]
    Destination port: (0..255) [-1]
    Basic stats (yes, y, no, n): [no]
    Extended stats (yes, y, no, n): [no]
    Trace reverse path (yes, y, no, n): [no]
    Source route (yes, y, no, n): [no]
    Timeout: (1..30) [10]


Target port is Embedded

Hop  In Port  Domain ID (Name)      Out Port    BW   Cost
------------------------------------------------------------
 0       E        10 (SW3900)          15       2G    500
 1       7         1 (swd3900TechPu     E       -      -
switch.admin>
```

The information that **pathInfo** provides is:

Hops	The number of switch-to-switch links (ISLs) transversed. The local switch is hop 0.
In Port	The port that the frames come in from on this path. For hop 0, the source port.
Domain ID	The domain ID of the switch
Name	The name of the switch
Out Port	The output port that the frames use to reach the next hop on this path. For the last hop, the destination port.
BW	The bandwidth of the output ISL, in Gbit/sec. It does not apply to the embedded port.
Cost	The cost of the ISL used by FSPF routing protocol. It applies only to an E_Port.

Using the FC-FC Routing Service

Chapter 11

The FC-FC (Fibre Channel) Routing Service provides Fibre Channel routing between two or more fabrics without merging those fabrics. The FC-FC Routing Service can be simultaneously used as a Fibre Channel router and for SAN extension over wide area networks (WANs) using FCIP.

The Fibre Channel Routing also supports interoperability with McDATA E/OS v7.x and 8.x. For more information about McDATA interoperability support, see "Connecting to McDATA SANs" on page 11-38.

The following sections describes the concepts and procedures for using Fibre Channel Routing Services and connecting to McDATA SANs.

- "Supported Platforms," next
- "Fibre Channel Routing Concepts" on page 11-2
- "Front Domain Consolidation" on page 11-5
- "Setting Up the FC-FC Routing Service" on page 11-11
- "FC Router Port Cost (Optional)" on page 11-26
- "Optional Configuration Procedures" on page 11-26
- "EX_Port Frame Trunking (Optional)" on page 11-30
- "Monitoring Resources" on page 11-33
- "Routing ECHO" on page 11-36
- "Upgrade and Downgrade Considerations" on page 11-37
- "Connecting to McDATA SANs" on page 11-38

Supported Platforms

FC-FC Routing is supported on the following platforms:

- SilkWorm FR4-18i blade
- Silkworm 7500 switch
- SilkWorm 48000 director when it is configured with an FR4-18i blade *and* uses chassis configuration option 5

Fibre Channel Routing Concepts

Fibre Channel routing introduces the following concepts:

- Logical storage area networks (LSANs)

  An LSAN is defined by zones in two or more edge or backbone fabrics that contain the same device(s). You can create LSANs that can span fabrics. These LSANs enable Fibre Channel zones to cross physical SAN boundaries without merging the fabrics while maintaining the access controls of zones (see Figure 11-1).

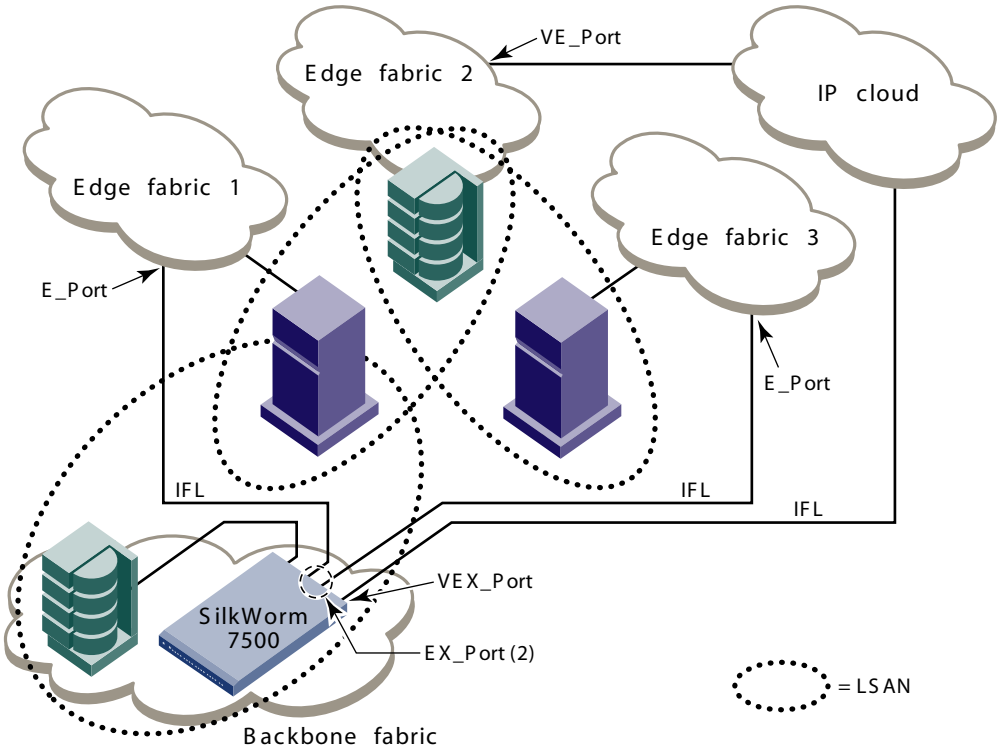

Figure 11-1 A metaSAN with Edge-to-Edge and Backbone Fabrics

Figure 11-1 shows a metaSAN with a backbone consisting of one SilkWorm 7500 connecting hosts in Edge Fabric 1 and 3 with storage in Edge Fabric 2 and the backbone through the use of LSANs. There are devices shared between the backbone and Edge Fabric 1, between Edge Fabric 1 and Edge Fabric 2, and between Edge Fabric 1 and Edge Fabric 2.

- EX_Port, VEX_Port

  Special types of ports, called an EX_Port and a VEX_Port, function similarly to an E_Port and VE_Port respectively, but terminate at the switch and do not propagate fabric services or routing topology information from one edge fabric to another. The link between an E_Port and EX_Port, or VE_Port and VEX_Port is called an *interfabric link* (IFL). You can configure multiple IFLs from a SilkWorm 7500, an FR4-18i operating in a SilkWorm 48000 using chassis configuration option 5, from additional Brocade routers, or from all three. These are referred to “an FC router” throughout this chapter and unless specified, any FC router can be used.

FC-FC Routing Services support EX_Ports and VEX_Ports connected to the same edge fabrics (see Figure 11-2). See Chapter 21, “Configuring and Monitoring FCIP Tunneling” for details about VE_Ports.

Figure 11-2 shows a metaSAN consisting of three edge fabrics connected through a SilkWorm 48000 containing an FR4-18i with interfabric links.

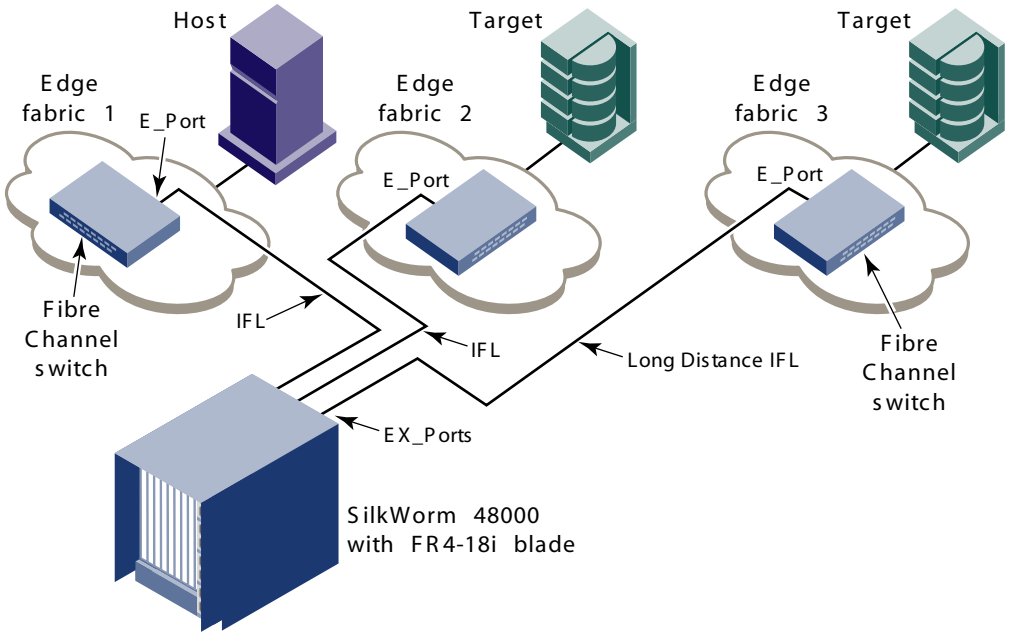

Figure 11-2 A metaSAN with Interfabric Links

- Edge Fabric

  A Fibre Channel fabric with targets and initiators connected through the supported platforms via an EX_Port or VEX_Port.

- Backbone Fabric

  A backbone fabric is an intermediate network that connects one or more edge fabrics. A backbone fabric consists of at least one FC router and possibly a number of Fabric OS-based Fibre Channel switches. It also enables hosts and targets in one edge fabric to communicate with devices in the other edge fabrics (see Figure 11-3). A backbone fabric also enables hosts and targets in one edge fabric to communicate with devices in other edge or backbone fabric.

  Note

  While the SilkWorm 7500 and SilkWorm 48000 with an FR4-18i blade facilitate communication between devices in edge fabrics with those in a backbone fabric, this is not true of the Multiprotocol Router AP7420.

- Fabric ID (FID)

  Every EX_Port and VEX_Port uses the FID property to identify the fabric at the opposite end of the IFL. Configure all of the EX_Ports and VEX_Ports attached to the same edge fabric with the same FID. The FID for every edge fabric must be unique from each backbone fabric's perspective. Configure the FID for all of the EX_Ports and VEX_Ports from a backbone fabric to uniquely reference each of the attached edge fabrics.

When two different backbones are connected to the same edge fabric, the backbone fabric IDs are different but the edge fabric FID should be the same. Configuring the same backbone fabric ID for two backbones that are connected to the same edge is invalid. In this configuration, a RAS log message is displayed warning about fabric ID overlap. Only two backbone fabrics not connected to the same edge can have the same backbone fabric ID.

- metaSAN

  A collection of SAN devices, switches, edge fabrics, Logical Storage Area Networks (LSANs), and SilkWorm Brocade Routers that comprise a physically connected but logically partitioned storage network. In this chapter, the term “metaSAN” is used for the collection of all SANs interconnected with Fibre Channel routers.

  A simple metaSAN can be constructed using a SilkWorm 7500, a SilkWorm 48000 director with an FR4-18i, or FC-FC to connect two or more separate fabrics. Additional FC routers can be used to increase the available bandwidth between fabrics and to provide redundancy.

- proxy PID

  The port ID of the proxy device. A proxy device is a virtual device presented into a fabric by a Fibre Channel router, and represents a real device on another fabric. When a proxy device is created in a fabric, the real Fibre Channel device is considered to be imported into this fabric. The presence of a proxy device is required for interfabric device communication. The proxy device appears to the fabric as a real Fibre Channel device, has a name server entry, and is assigned a valid port ID. The port ID is only relevant on the fabric in which the proxy device has been created.

- phantom domains

  The Fibre Channel router emulates two levels of phantom domains. The first set of phantom domains are known as “front phantom domains.” There is one front phantom domain from FCR to an edge fabric.

  The second level of phantom domains is known as a “translate phantom domain.” The EX_Ports also present translate phantom domains in edge fabrics as being topologically behind the front domains; if the translate phantom domain is in a backbone fabric, then it is topologically present behind the Fibre Channel router because there is no front domain in a backbone fabric.The translate phantom domain is a router virtual domain that represents an entire fabric. Device connectivity can be achieved from one fabric to another—over the backbone or edge fabric through this virtual domain—without merging the two fabrics. Translate phantom domains are sometimes referred to as “translate domains,” or “xlate domains.” If an FR4-18i blade is attached to an edge fabric using an EX_Port, it will create translate phantom domains in the fabric corresponding to the imported edge fabrics with active LSANs defined. If you import devices into the backbone fabric, then a translate phantom domain is created in the backbone device (in addition to the one in the edge fabric).

  If you lose connectivity to the edge fabric due to link failures or the IFL being disabled (for example, by issuing the **portDisable** command from the switch in the attached fabric), translate phantom domains remain visible. This prevents unnecessary fabric disruptions caused by translate phantom domains repeatedly going offline and online due to corresponding IFL failures. To remove the translate phantom domain in the backbone, disable all EX_Ports or VEX_Ports through which the translate phantom domain was created.

Figure 11-3 shows another metaSAN consisting of a host in Edge SAN 1 connected to storage in Edge SAN 2 through a backbone fabric connecting two SilkWorm 48000s, each containing FR4-18is.

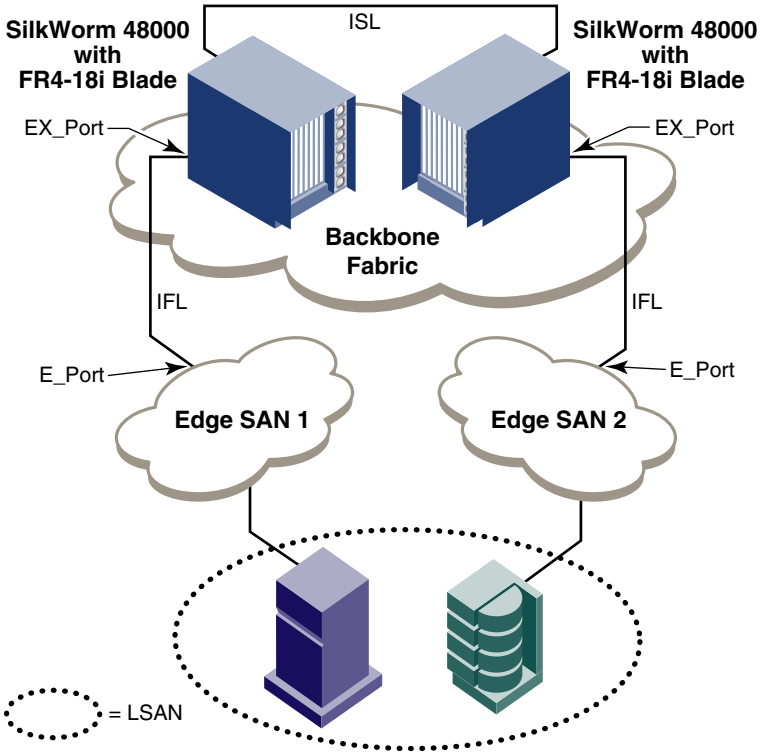

Figure 11-3 Edge SANs Connected Through a Backbone Fabric

Front Domain Consolidation

Fabric OS v5.2.0 provides Front Domain Consolidation. With this feature, one front domain from a single FCR is projected to the edge fabric regardless of the number of EX_Ports connected from that router to the edge fabric. The front domain is consolidated only within the FCR switch. Another router connected to the same edge fabric projects a different front domain.

The FCR switch is connected to the edge fabrics via EX_Ports or VEX_Ports. On prior FCR releases (XPath v7.1.2, v7.3, v7.4 and Fabric OS v5.1) every EX_Port connecting to an edge fabric projects its own front domain. For example, if there are three EX_Ports connected from the router to the edge fabric, there will be three front domains projected to the edge fabric.

Front domain consolidation conserves resources on the FCR switch. Expanded use of FCR switches with multiple EX_Ports connected to the same edge fabric raises the overhead of projecting a front domain for each EX_Port and quickly uses up resources within the edge fabric and resource consumption on the FCR.

In previous releases, every EX_Port connected to the same edge fabric had a unique domain and unique domains require separate WWNs. With front domain consolidation, the domain is same for all ports connected to same edge fabric from the FC-FC Router, so in this case the WWN will be the same.

Supported Configurations and Platforms

In an edge fabric that contains a mix of AD-capable switches and switches that are not aware of AD, FCR should be connected directly to the AD-capable switch. For more information, see "Use of Administrative Domains with LSAN Zones and FCR" on page 11-22.

The supported configurations are:

- FCR connected to a Brocade nonsecured fabric
- FCR connected to a Brocade secured fabric
- FCR connected to a McDATA Open mode
- FCR connected to a McDATA Fabric mode
- FCR connected to a Brocade secured and non secured fabrics with EX_Port trunking enabled
- Can interoperate with older FCR switches (XPath v7.4.x and Fabric OS v5.1)

Upgrade and Downgrade Considerations

The following considerations apply when upgrading to or downgrading from Fabric OS 5.2.0 with front domain consolidation:

- During an upgrade to Fabric OS v5.2 from Fabric OS v5.1:
  - The router switch is changed from one front domain per EX_Port to a shared front domain for the EX_Ports that are connected to the same edge fabric.
  - One port per edge fabric remains online and connected to the edge fabric. The other ports connected to the same edge fabric are reinitialized (port offline and online). The ports share the front domain of the port that remained online.
  - Minimal traffic disruption occurs when you have more than one IFL connection to the edge fabric.
  - No traffic disruption should occur in the case of a single IFL to an edge fabric.
  - Frame drops can occur as routing changes from multiple ports to one port.
- During a downgrade from Fabric OS v5.2 to a prior Fabric OS version (for example, Fabric OS v5.1):
  - The router switch is changed from a consolidated front domain to multiple front domains and all the EX_Ports will be reinitialized. This is a disruptive procedure.
  - Traffic disruption can occur.

Using Front Domain Consolidation

To support the consolidated front domain feature, existing CLI and API commands are used to manage and configure the router. Existing CLI commands offer additional options to support the front domain consolidation feature.

The **portCfgExport** command has additional options to verify the front domain ID. The **portCfgExport –d** option is changed to enforce use of the same front domain ID for the EX_Ports connected to the same edge fabric. The **portCfgExport** display results remain the same.

For more information about the **portCfgExport -d** option, see "portCfgExport Options," on page 11-18 and the command details in the *Fabric OS Command Reference*.

The following example illustrates the use of the **portcfgexport** command.

```
switch:admin_06> portcfgexport 2/0
Port 2/0 info
Admin: enabled
State: OK
Pid format: core(N)
Operate mode: Brocade Native
Edge Fabric ID: 16
Front Domain ID: 160
Front WWN: 50:06:06:9e:20:9f:ce:10
Principal Switch: 7
principal WWN: 10:00:00:60:69:c0:05:8a
Fabric Parameters: Auto Negotiate
R_A_TOV: 9000(N)
E_D_TOV: 2000(N)
Edge fabric's primary wwn: N/A
Edge fabric's version stamp: N/A
```

The **fabricShow** command shows only one front domain per a router connection. In the following example **fcr_fd_3_32** has two EX_Port connections but one front domain ID.

```
switch:admin_06> fabricshow
Switch ID   Worldwide NameEnet IP Addr    FC IP AddrName
-------------------------------------------------------------------------
1: fffc01 50:00:51:e3:70:d2:5f:cb  0.0.0.00.0.0.0"fcr_xd_1_32"
2: fffc02 50:00:51:e3:70:d2:5f:cc  0.0.0.00.0.0.0"fcr_xd_2_1"
3: fffc03 50:00:51:e3:60:14:0e:20  0.0.0.00.0.0.0"fcr_fd_3_32"
199: fffcc7 10:00:00:60:69:c0:05:d1  10.32.156.33    0.0.0.0
>"mojo__10" The Fabric has 4 switches
```

For more information about the **fabricShow** command, see the *Fabric OS Command Reference*.

Range of Output Ports

The edge fabric detects only one front domain from an FC router connected through multiple output ports. The output port of the front domain is not fixed to 0; the values can be in a range of 129–255. The range of the output ports connected to the xlate domain is also 129–255. This range enables the front domain to connect to 127 remote xlate domains.

To display the range of output ports connected to the xlate domains

1. Log in to the FC router.
2. Enter the **lsDbShow** command on the edge fabric.

   The following example shows the range of output ports.

```
linkCnt = 2,    flags = 0x0
LinkId =  53, out port =   1, rem port =  35, cost = 500, costCnt = 0, type = 1
LinkId =  57, out port =  129, rem port =  18, cost = 500, costCnt = 0, type = 1
```

The following example also shows the use of the **lsDbShow** display on the edge fabric. The front domain, domain 3, has two links representing two EX_Port connections with output ports 129 and 132.

```
Domain = 3, Link State Database Entry pointer = 0x100bbcc0
.........
linkCnt = 4,  flags = 0x0
LinkId = 199, out port = 129, rem port =   2, cost = 10000, costCnt = 0, type = 1
LinkId = 199, out port = 132, rem port =   3, cost = 10000, costCnt = 0, type = 1
LinkId =   2, out port =   1, rem port =   2, cost = 10000, costCnt = 0, type = 1
LinkId =   1, out port =  32, rem port =   2, cost = 10000, costCnt = 0, type = 1
```

To verify normal operation

1. Connect the FC router to an edge fabric switch via multiple EX_Ports from the same router.
2. Confirm that the ports are enabled with **portCfgShow**.
3. Confirm that the EX_Ports share the same front domain PID and node WWN, using the **portCfgExport** command.

Note

Abnormal operation for front domain consolidation exists when the ports do not come online or in the Normal Operation setup described previously, the EX_Ports do not share the same front domain consolidation PID and node WWN.

Support

In case of an unexpected failure at the customer site, save the output from the **supportShow** and **supportSave** run on the FC router and edge fabric switches.

Proxy Devices

A SilkWorm 7500 or 48000 with an FR4-18i achieves interfabric device connectivity by creating proxy devices (hosts and targets) in attached fabrics that represent real devices in other fabrics. For example, a host in Fabric 1 can communicate with a target in Fabric 2 as follows:

- A proxy target in Fabric 1 represents the real target in Fabric 2.
- Likewise, a proxy host in Fabric 2 represents the real host in Fabric 1.

The host discovers and sends Fibre Channel frames to the proxy target. The SilkWorm 7500 or 48000 with an FR4-18i receives these frames, translates them appropriately, then delivers them to the destination fabric for delivery to the target.

The target responds by sending frames to the proxy host. Hosts and targets are exported from the edge SAN to which they are attached and, correspondingly, imported into the edge SAN reached through Fibre Channel routing. Figure 11-4 illustrates this concept.

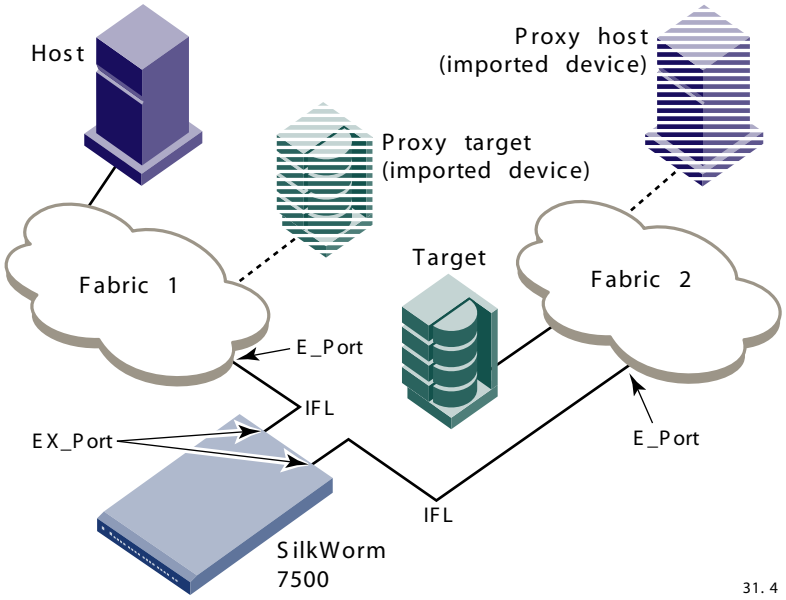

Figure 11-4 metaSAN with Imported Devices

Routing Types

- Edge-to-Edge

  Occurs when devices in one edge fabric communicate with devices in another edge fabric through one or more Fibre Channel routers.

- Backbone-to-Edge

  Fibre Channel routers can connect to a common fabric—known as a backbone fabric—via E_Ports. A backbone fabric can be used as a transport fabric that interconnects edge fabrics. Fibre Channel routers also enable hosts and targets in edge fabrics to communicate with devices in the backbone fabric, known as “backbone-to-edge routing.” From the edge fabric's perspective, the backbone fabric is just like any other edge fabric. For the edge fabric and backbone fabric devices to communicate, the shared devices need to be presented to each other's native fabric.

  To do so, at least one translate phantom domain (switch) is projected into the backbone fabric. This translate phantom switch represents the entire edge fabric. The shared physical device in the edge has a corresponding proxy device on the translate phantom domain switch.

Each edge fabric has one and only one xlate switch to the backbone fabric. The backbone fabric device communicates with the proxy devices whenever it needs to contact the shared physical device in the edge. The FC-FC Routing Service receives the frames from the backbone switches destined to the proxy device, and redirects the frame to the actual physical device. As with an edge fabric, the translate phantom domain switch can never be the principal switch of the backbone fabric. Front domains are not created; rather, only translate phantom domains are created in the backbone fabric.

Devices are exported from the backbone fabric to one or more edge fabrics using LSANs. See "Configuring LSANs and Zoning" on page 11-22 for more information.

Note

Secure Fabric OS, Management Server Platform services, and interopmode are not supported in the backbone fabric.

Fibre Channel NAT and Phantom Domains

Within an edge fabric or across a backbone fabric, the standard Fibre Channel FSPF protocol determines how frames are routed from the source Fibre Channel (FC) device to the destination FC device. The source or destination device can be a proxy device. When frames traverse the fabric through a SilkWorm 7500 or 48000 in the backbone (BB), the frames are routed to another EX_Port or VEX_Port.

Fibre Channel fabrics require that all ports (EX_Ports or VEX_Ports) be identified by a unique PID. In a single fabric, FC protocol guarantees that domain IDs are unique, and so a PID formed by a domain ID and area ID is unique within a fabric. However, the domain IDs and PIDs in one fabric might be duplicated within another fabric, just as IP addresses unique to one private network are likely to be duplicated within another private network.

In an IP network, a network router can maintain network address translation (NAT) tables to replace private network addresses with public addresses when a packet is routed out of the private network, and replace public addresses with private addresses when a packet is routed from the public network to the private network. The Fibre Channel routing equivalent to this IP-NAT is the Fibre Channel network address translation (FC-NAT). Using FC-NAT, the proxy devices in a fabric can have different PIDs than the real devices that they represent, allowing the proxy devices to have appropriate PIDs for the address space of their corresponding fabric.

All EX_Ports connected to same edge fabric from one physical FC router present a single front phantom domain and one additional translate (xlate) phantom domain for each edge fabric accessed through it. All EX_Ports and VEX_Ports connected to an edge fabric use the same xlate domain ID number for an imported edge fabric; this value persists across switch reboots and fabric reconfigurations. Xlate domains are presented as being connected topologically behind one or more front domains; each FC-Router presents one front domain to edge fabric. This allows redundant paths to remote fabrics to present redundant paths to proxy devices to an edge fabric.

Phantom domains are like logical switches that appear to be connected to an edge fabric through the front domains that are presented by EX_Ports and/or VEX_Ports. The combination of front domains and xlate domains allows routing around path failures, including path failures through the routers. The multiple paths to an xlate domain provide additional bandwidth and redundancy.

There are some differences in how the xlate domain is presented in the BB. The BB xlate domains are topologically connected to FC routers and participate in FC-Protocol in BB. Front domains are not needed in the BB. As in the case of xlate domain in edge fabric, BB xlate domains provide additional bandwidth and redundancy by being able to present itself being connected to single or multiple FC routers with each FC router capable of connecting multiple IFLs to edge fabrics.

Setting Up the FC-FC Routing Service

To set up the FC-FC Routing Service, perform the following tasks:

1. "Performing Verification Checks" next
2. "Assigning Backbone Fabric IDs" on page 11-13
3. "Configuring FCIP Tunnels (Optional)" on page 11-14
4. "Configuring FC-FC Routing to Work with Secure Fabric OS (Optional)" on page 11-14
5. "Configuring an Interfabric Link" on page 11-16
6. "FC Router Port Cost (Optional)" on page 11-26
7. "EX_Port Frame Trunking (Optional)" on page 11-30
8. "Configuring LSANs and Zoning" on page 11-22

See Chapter 9, "Configuring SilkWorm Directors" for more details about configuration options for Brocade directors.

Performing Verification Checks

Before configuring a fabric to connect to another fabric, you must perform the following verification checks on the switch or director.

To perform verification checks

1. Log in to the switch or director as admin and enter the **version** command. Verify that Fabric OS v5.2.0 is installed on the SilkWorm 7500 or FR4-18i blade, as shown in the following example.

```
switch:admin_06> version
Kernel:     2.4.19
Fabric OS:  v5.2.0
Made on:    Thu Sep 21 01:15:34 2006
Flash:      Fri Sep 22 20:53:48 2006
BootProm:   4.5.3
```

2. If configuring the SilkWorm 48000 with an FR4-18i blade, enter the **slotShow** command to verify that the FR4-18i blade is present.

```
switch:admin_06> slotshow
Slot   Blade Type   ID    Status
---------------------------------
  1     SW BLADE   17     ENABLED
  2     AP BLADE   24     ENABLED
  3     UNKNOWN           VACANT
  4     UNKNOWN           VACANT
  5     CP BLADE   16     ENABLED
  6     CP BLADE   16     ENABLED
  7     SW BLADE   17     ENABLED
  8     UNKNOWN           VACANT
  9     SW BLADE   18     ENABLED
 10     UNKNOWN           VACANT
```

See Chapter 9, “Configuring SilkWorm Directors” for a list of blades and their corresponding IDs.

3. If configuring the SilkWorm 48000 with an FR4-18i blade, then enter the **chassisConfig** command to verify that the director is using configuration option 5.

```
switch:admin_06> chassisconfig

Current Option: 5

All Supported Options
-------------------------------------------------
Option 1: One 128-port switch
    Blade ID's 4, 17 in slots 1-4, 7-10
    Blade ID's 5, 16 in slots 5-6
Option 5: One 256-port switch
    Blade ID's 17, 18, 24 in slots 1-4, 7-10
    Blade ID 16 in slots 5-6

Please use slotshow to see Blade IDs currently in the system.
```

4. Enter the **interopMode** command and verify that SilkWorm switch interoperability with switches from other manufacturers is disabled.

```
switch:admin_06> interopmode
InteropMode: Off

Usage: InteropMode 0|1
         0: to turn it off
         1: to turn it on
```

5. Enter the **secModeShow** command to verify that security is disabled.

```
switch:admin_06> secmodeshow
Secure Mode: DISABLED.
```

6. Enter the **msPlatShow** command to verify that Management Server Platform database is disabled in the backbone fabric.

```
switch:admin_06> msplatshow
*MS Platform Management Service is NOT enabled.
```

If any of the items listed in the prior steps are enabled, you can see the *Fabric OS Command Reference* for information on how to disable the option. For information about security, see Chapter 4, “Configuring Standard Security Features” and Chapter 5, “Configuring Advanced Security”.

When it is in strict mode, ACL cannot support Fibre Channel routing in the fabric. Before connecting an edge fabric to an FC router and before setting up the FC router in the BB, verify that the Fabric Wide Consistency Policy is not in ‘strict’ mode by issuing the **fddCfg --showall** command.

If the Fabric Wide Consistency Policy has the ‘S’ letter in it in the edge fabric or the BB fabric, do not connect the edge fabric or the BB to the FC router. See the *Fabric OS Command Reference* for details.

Assigning Backbone Fabric IDs

If your configuration has only one backbone fabric, then this task is not required because the backbone fabric ID in this situation defaults to a value of 1.

All switches in a backbone fabric must have the same backbone fabric ID. You can configure the backbone fabric ID using the **fcrConfigure** command. The backbone fabric ID is required to be unique from the perspective of every attached edge fabric. Fabric ID changes made on a switch are not propagated to other switches in backbone fabric. Rather, the backbone fabric administrator is responsible for making sure that all switches in the backbone have the same fabric ID. Because fabric IDs are used heavily by the routing protocol between the Fibre Channel routers, using the wrong fabric ID can affect both edge-to-edge and backbone-to-edge routing.

In addition to ensuring that the backbone fabric IDs are the same within the same backbone, you need to make sure that when two different backbones are connected to the same edge fabric, the backbone fabric IDs are different (but the edge fabric FID should be the same). Configuration of two backbones with the same backbone fabric ID that are connected to the same edge is invalid. In this configuration, a RAS log message is displayed warning about fabric ID overlap. However, when two backbone fabrics are not connected to the same edge, they can have the same backbone fabric ID.

To assign backbone fabric IDs

1. Log in to the switch or director.
2. Enter the **fosConfig --disable fcr** command to disable the FC-FC Routing Service. See the *Fabric OS Command Reference* or the CLI man pages for more information about the **fosConfig** command.

Note

The default state for the FCR is disabled. The **fcrEnable** and **fcrDisable** commands continue to operate as before, however **fosConfig** is the preferred method for enabling and disabling the FC-FC Routing (FCR) Service.

3. Enter the **fcrConfigure** command. At the prompt, enter the fabric ID, or press **Enter** to specify the default fabric ID (1).
4. Verify that the fabric ID is different from that set for edge fabrics.

   Multiple FC routers attached to the same backbone fabric must have the same backbone fabric ID.
5. Then enter the **fosConfig --enable fcr** command.

```
switch:admin_06> fosconfig --disable fcr
FC Router service is disabled

switch:admin_06> fcrconfigure
FC Router parameter set. <cr> to skip a parameter
Backbone fabric ID: (1-128)[1]

switch:admin_06> fosconfig --enable fcr
FC Router service is enabled
```

Configuring FCIP Tunnels (Optional)

The optional Fibre Channel over IP (FCIP) Tunneling Service enables you to use “tunnels” to connect instances of Fibre Channel SANs over IP-based networks to transport all Fibre Channel ISL and IFL traffic. FCIP is a prerequisite for configuring VEX_Ports; if you are only using FC_Ports, then there is no need to perform this step.

Note

If using FCIP in your FC-FC Routing configuration, you must first configure FCIP tunnels. Once a tunnel is created, it defaults to a disabled state. Then configure the VE_Port or VEX_Port. After the appropriate ports are configured, enable the tunnel.

See Chapter 21, “Configuring and Monitoring FCIP Tunneling” for instructions on how to configure FCIP tunnels.

Configuring FC-FC Routing to Work with Secure Fabric OS (Optional)

If you do not have Secure Fabric OS enabled in the edge fabric, then you are not required to complete the tasks in this section.

Note

Secure Fabric OS is *not* supported in backbone fabrics.

The SilkWorm 7500 and SilkWorm 48000 with an FR4-18i blade support Fibre Channel routing between Brocade’s secure fabric employing Secure Fabric OS via DH-CHAP (Diffie-Hellman with Challenge-Handshake Authentication Protocol) authentication. It also supports secure fabric to nonsecure fabrics. Secure Fabric OS is an optionally licensed product that provides customizable security restrictions through local and remote management channels on a SilkWorm fabric.

The FC-FC Routing Service uses only the DH-CHAP shared secrets to provide switch-to-switch authentication when connecting to a Secure Fabric OS fabric.

To determine whether or not an EX_Port or VEX_Port is connected to a Secure Fabric OS fabric, enter the **portShow**, **portCfgEXPort,** or **portCfgVEXPort** command, as described in the *Fabric OS Command Reference*. Note that you should issue these commands only after the IFLs have been configured and for VEX ports, the FCIP tunnel(s) are up and running. For more details, see “Configuring an Interfabric Link” on page 11-16 and “Configuring FCIP Tunnels (Optional)” on page 11-14.

Configuring Secure Fabric OS DH-CHAP Secret

While Secure Fabric OS supports the SLAP, FCAP and DH-CHAP authentication protocols to communicate with each switch, Fabric OS v5.2.0 (and Fabric OS v5.1.0) supports only DH-CHAP.

The SilkWorm 7500 and 48000 with an FR4-18i do not initiate DH-CHAP authentication requests; rather, they respond to DH-CHAP requests only from the edge switch to which they are connected—in this case, the Secure Fabric OS switch.

As soon as you connect the SilkWorm 7500 or 48000 with an FR4-18i blade to a Secure Fabric OS switch, DH-CHAP authentication is initiated.

The DH-CHAP secrets are configured both on the Secure Fabric OS switch and the SilkWorm 7500 or 48000 with an FR4-18i blade. Each entry specifies the WWN of the peer to which it is connected. For example, on the SilkWorm 48000 with an FR4-18i blade, specify the WWN of the Secure Fabric OS switch and the secrets. On the Secure Fabric OS switch, specify the WWN of the front domain (EX_Port or VEX_Port) and the secrets. To view the front domain WWN, issue the **portCfgEXPort** command on the Fibre Channel router side.

The WWN of the front domain (EX_Port or VEX_Port) that is connected to the Secure Fabric OS switch should be present in the Switch Connection Controls (SCC) list. See the *Secure Fabric OS Administrator's Guide* for details about the SCC and other Secure Fabric OS features.

To configure a DH-CHAP secret word

1. Log in to the SilkWorm 7500 or 48000 with an FR4-18i blade with administrative privileges.
2. At the telnet prompt, enter the **secAuthSecret** command. The secret must be between 8 and 40 characters long.

   Setting up secret keys does not initiate DH-CHAP authentication. DH-CHAP authentication is performed whenever a port or a switch is enabled.
3. Follow the instructions provided on screen, as shown in the following example.
   a. Type the port or switch WWN.

   Note
   Use only the WWN as the input. The domain ID or switch name is not acceptable.

   b. Type and confirm the peer secret.
   c. Type and confirm the local secret.
4. After you have added all of the DH-CHAP secret information, press **Enter** to indicate that you have completed the secret key setup.

5. When prompted, type **y**. The DH-CHAP secret is now stored in the secret word database and is ready for use.

```
switch:admin_06> secauthsecret --set
This command is used to set up secret keys for the DH-CHAP authentication.
The minimum length of a secret key is 8 characters and maximum 40
characters. Setting up secret keys does not initiate DH-CHAP
authentication. If switch is configured to do DH-CHAP, it is performed
whenever a port or a switch is enabled.

Warning: Please use a secure channel for setting secrets. Using
an insecure channel is not safe and may compromise secrets.

Following inputs should be specified for each entry.

1. WWN for which secret is being set up.
2. Peer secret: The secret of the peer that authenticates to peer.
3. Local secret: The local secret that authenticates peer.

Press enter to start setting up secrets >
Enter WWN, Domain, or switch name (Leave blank when done):
10:00:00:60:69:80:05:14
Enter peer secret:
Re-enter peer secret:
Enter local secret:
Re-enter local secret:
Enter WWN, Domain, or switch name (Leave blank when done):
Are you done? (yes, y, no, n): [no] y
Saving data to key store... Done.
```

To view a DH-CHAP secret word database

1. Log in as admin to the SilkWorm 7500 or 48000 with an FR4-18i blade.
2. At the telnet prompt, enter the **secAuthSecret** command as shown:

```
switch:admin_06> secauthsecret --show

WWN                              DId                       Name
-----------------------------------------------------------------
10:00:00:60:69:80:05:14          1                         switch
```

For details about the **setAuthSecret** command, see the *Secure Fabric OS Administrator's Guide*.

Configuring an Interfabric Link

Before configuring an IFL, be aware that you cannot configure both IFLs (EX_Ports, VEX_Ports) and ISLs (E_Ports) from a single SilkWorm 7500 or FR4-18i blade to the same edge fabric.

Configuring an interfabric link involves disabling ports and cabling them to other fabrics, configuring those ports for their intended use, and then enabling the ports.

When it is in strict mode, ACL cannot support Fibre Channel routing in the fabric. Before connecting an edge fabric to an FC router and before setting up the FC router in the BB, verify that the Fabric Wide Consistency Policy is not in 'strict' mode by issuing the **fddCfg --showall** command.

If the Fabric Wide Consistency Policy has the 'S' letter in it in the edge fabric or the BB fabric, do not connect the edge fabric or the BB to the FC router.

Note

To ensure that fabrics remain isolated, disable the port prior to inserting the cable. If you are configuring an EX_Port, disable the port prior to making the connection.

To configure an IFL for both edge and backbone connections

1. On the SilkWorm 7500 or FR4-18i blade, disable the port that you are configuring an EX_Port (the one connected to the SilkWorm switch) by issuing the **portDisable** command.

```
switch:admin_06> portdisable 7/10
```

   You can verify that port 7 has been disabled by issuing the **portShow** command for the port.

2. Configure each port that connects to an edge fabric as an EX_Port or VEX_Port. Note the following:

   - **portCfgVEXPort** works only on the GbE ports.
   - **portCfgEXPort** (only on the FC ports on the SilkWorm 7500 or SilkWorm 48000 with an FR4-18i blade) commands work only on ports that are capable of FC-FC routing.

   Use the **portCfgEXPort** or **portCfgVEXPort** command to:

   - Enable or disable EX_Port or VEX_Port mode.
   - Set the fabric ID (avoid using fabric ID 1, which is the default for backbone connections).

   The following example enables the EX_Port (or VEX_Port) and assigns a Fabric ID of 30 to port 7.

```
switch:admin_06> portcfgexport 7 -a 1 -f 30
switch:admin_06> portcfgexport 7/10
Port   7/10   info
Admin:                  enabled
State:                  NOT OK
Pid format:             Not Applicable
Operate mode:           Brocade Native
Edge Fabric ID:         30
Preferred Domain ID:    160
Front WWN:              50:06:06:9e:20:38:6e:1e
Fabric Parameters:      Auto Negotiate
R_A_TOV:                Not Applicable
E_D_TOV:                Not Applicable
Authentication Type: None
DH Group: N/A
Hash Algorithm: N/A
Edge fabric's primary wwn: N/A
Edge fabric's version stamp: N/A
```

portCfgExport Options

This port can now connect to another switch. The following list describes the options for the **portCfgExport** command. For more information about the **portCfgExport** and **portCfgVexport** commands, see the *Fabric OS Command Reference*.

-a Sets the EX_Port to enabled (1) or disabled (2). Admin use only.

-f Sets the fabric ID (1 to 128). Each edge fabric must have a unique ID, and EX_Ports (or VEX_Ports) connected to the same edge fabric must have the same fabric ID. The default value is the port number divided by 3, plus 2 and rounded down.

-r R_A_TOV used for port negotiation (2000 - 120000).

-e E_D_TOV used for port negotiation (1000 - 60000).

-d Preferred domain ID (1-239). This command enforces the use of the same preferred domain ID for all the ports connected to the same edge fabric.

When this option is specified, the preferred domain ID is compared against the online ports. If the domain ID are different, an error message is issued and the command fails.

When the -d option is not specified, if there are online ports connected to the same edge fabric, the preferred domain ID is set to the preferred domain ID of those online ports. Otherwise, if there are offline ports that are set to EX-Port, the preferred domain ID is set to those offline ports. If none of the above conditions apply, the existing value is left untouched.

-p PID format (0-native, 1-core, 2-extended-edge). The value must match the edge fabric setting. The default value is 1.

-t Negotiate fabric parameters (1-enable, 2-disable)

-m Sets the port mode. (0-Brocade, 1-Open, 2-McDATA Fabric, 3-McDATA Fabric Legacy). This command enforces the use of the same port mode for all the ports connected to the same edge fabric.

When this option is specified, the port mode is compared against the online ports. If the modes are different, an error message is issued and the command fails.

When the -m option is not specified:

If there are online ports connected to the same edge fabric, the mode is set to the mode of those online ports. Otherwise, if there are offline ports that are set to EX_Port, the mode is set to those offline ports.

If none of the above conditions apply, the existing value is left untouched.

For related FC-FC Routing commands, seesee **fcrxlateconfig**, **fcrconfigure**, and **fcrproxyconfig** in the *Fabric OS Command Reference*.

A Fibre Channel router can interconnect multiple fabrics. EX_Ports or VEX_Ports attached to more than one edge fabric must configure a different fabric ID for each edge fabric.

At this point you have some options to consider before proceeding to the next step. These options include FCR router port cost operations and setting up either ISL or EX_Port trunking. For information about using FCR Router Port Cost operations, see “FC Router Port Cost (Optional)” on page 11-26 and for information on trunking setup, see “Using EX_Port Frame Trunking” on page 11-31.

3. Enter the **portEnable** command to enable the ports that you disabled in step 2. You can now physically attach ISLs from the Fibre Channel Router to the edge fabric.

```
switch:admin_06> portenable 7/10
```

4. Enter the **portCfgShow** command to view ports that are persistently disabled.

```
switch:admin_06> portcfgshow 7/10
Area Number:              74
Speed Level:              AUTO
Trunk Port                OFF
Long Distance             OFF
VC Link Init              OFF
Locked L_Port             OFF
Locked G_Port             OFF
Disabled E_Port           OFF
ISL R_RDY Mode            OFF
RSCN Suppressed           OFF
Persistent Disable        OFF
NPIV capability           ON
EX Port                   ON
Mirror Port               ON
```

5. After identifying such ports, enter the **portCfgPersistentEnable** command to enable the port, and then the **portCfgShow** command to verify the port is enabled.

```
switch:admin_06> portcfgpersistentenable 7/10

switch:admin_06> portcfgshow 7/10
Area Number:              74
Speed Level:              AUTO
Trunk Port                OFF
Long Distance             OFF
VC Link Init              OFF
Locked L_Port             OFF
Locked G_Port             OFF
Disabled E_Port           OFF
ISL R_RDY Mode            OFF
RSCN Suppressed           OFF
Persistent Disable        OFF
NPIV capability           ON
EX Port                   ON
Mirror Port               ON
```

6. Enter either the **portCfgEXPort** or **portShow** command to verify that each port is configured correctly:

```
switch:admin_06> portcfgexport 7/10

        Port   7/10   info
Admin:                  enabled
State:                  NOT OK
Pid format:             Not Applicable
Operate mode:           Brocade Native
Edge Fabric ID:         30
Preferred Domain ID:    160
Front WWN:              50:06:06:9e:20:38:6e:1e
Fabric Parameters:      Auto Negotiate
R_A_TOV:                Not Applicable
E_D_TOV:                Not Applicable
Authentication Type: None
DH Group: N/A
Hash Algorithm: N/A
Edge fabric's primary wwn: N/A
Edge fabric's version stamp: N/A
```

```
switch:admin_06> portshow 7/10
portName:
portHealth: OFFLINE

Authentication: None

EX_Port Mode:   Enabled
Fabric ID:      30
Front Phantom:  state = Not OK  Pref Dom ID: 160
Fabric params:  R_A_TOV: 0    E_D_TOV: 0    PID fmt: au
to

Authentication Type: None
Hash Algorithm: N/A
DH Group: N/A
Edge fabric's primary wwn: N/A
Edge fabric's version stamp: N/A

portDisableReason: None
portCFlags: 0x1
portFlags: 0x1   PRESENT U_PORT EX_PORT
portType:  10.0
portState: 2    Offline
portPhys:  2    No_Module
portScn:   0
port generation number:    0
portId:    014a00
portIfId:    4372080f
portWwn:   20:4a:00:60:69:e2:03:86
portWwn of device(s) connected:

Distance:  normal
portSpeed: N4Gbps

LE domain: 0
Interrupts:        0          Link_failure: 0          Frjt :         0
Unknown:           0          Loss_of_sync: 0          Fbsy :         0
Lli:               0          Loss_of_sig:  2
Proc_rqrd:         0          Protocol_err: 0
Timed_out:         0          Invalid_word: 0
Rx_flushed:        0          Invalid_crc:  0
Tx_unavail:        0          Delim_err:    0
Free_buffer:       0          Address_err:  0
Overrun:           0          Lr_in:        0
Suspended:         0          Lr_out:       0
Parity_err:        0          Ols_in:       0
2_parity_err:      0          Ols_out:      0
CMI_bus_err:       0

Port part of other ADs: No
```

7. Enter the **switchShow** command to verify the EX_Port (or VEX_Port), edge fabric ID, and name of the edge fabric switch (containing the E_Port or VE_Port).
8. Enter the **fcrFabricShow** command to view any edge fabric's switch names and ensure links are working as expected:

```
switch:admin_06> fcrfabricshow
FC Router WWN: 10:00:00:60:69:e2:03:86, Dom ID:   1, Info: 10.33.59.38,
"fcr_meteor3"
   EX_Port       FID    Neighbor Switch Info (WWN, enet IP, name)
   ------------------------------------------------------------------------
      74          30     10:00:00:60:69:e2:09:fa   10.33.59.25 "fcr_meteor2"
```

Configuring LSANs and Zoning

An LSAN consists of zones in two or more edge or backbone fabrics that contain the same device(s). LSANs essentially provide selective device connectivity between fabrics without forcing you to merge those fabrics. FC routers provide multiple mechanisms to manage interfabric device connectivity through extensions to existing switch management interfaces. You can define and manage LSANs using Brocade Advanced Zoning or Brocade Fabric Manager.

Use of Administrative Domains with LSAN Zones and FCR

You can create LSAN zones as a physical fabric administrator or as an individual administrative domain (AD) administrator. The LSAN zone can be part of root zone database or the AD zone database. FCR harvests the LSAN zones from all administrative domains. If both edge fabrics have the matching LSAN zones and both devices are online, FCR triggers a device import. To support legacy applications, WWNs are reported based on the administrative domain context. As a result, you must not use the NAA field in the WWN to detect an FC Router. LSAN zone enforcement in the local fabric occurs only if the administration domain member list contains both of the devices (local and imported device) specified in the LSAN zone.

For more information, see Chapter 7, "Managing Administrative Domains".

Defining and Naming Zones

Zones are defined locally. Names and memberships, with the exception of hosts and targets exported from one fabric to another, do not need to be coordinated with other fabrics. For example, in Figure 11-3, when the zones for Edge SAN 1 are defined, you do not need to consider the zones in Edge SAN 2, and vice versa.

Zones that contain hosts and targets that are shared between the two fabrics need to be explicitly coordinated. Although an LSAN is managed using the same tools as any other zone on the edge fabric, two behaviors distinguish an LSAN from a conventional zone:

- A required naming convention. The name of an LSAN begins with the prefix "LSAN_". The LSAN name is letter case insensitive; for example, *lsan_* is equivalent to *LSAN_*, *Lsan_*, and so on.
- Members must be identified by their port WWN because PIDs are not necessarily unique across fabrics. The names of the zones need not be explicitly the same, and membership lists of the zones need not be in the same order.

LSAN Zones and Fabric-to-Fabric Communications

Because zoning is enforced by all involved fabrics, any communication from one fabric to another must be allowed by the zoning setup on both fabrics. If the SANs are under separate administrative control, then separate administrators maintain access control.

If you are managing other switches in a fabric, it is recommended that you run the **defZone --show** command on your Fabric OS v5.1.0 or later switches as a precaution. Default zoning behavior in Fabric OS v5.1.0 and later operates differently compared to other Fabric OS versions (versions 2.x, 3.x 4.x and 5.0.1).

For example, if you issue the **defZone --noaccess** command on a Fabric OS v5.1.0 or later switch, then default zoning configurations will be created on each switch in the fabric (v2.x, v3.x, v4.x or v5.0.1 switches). Fabric OS v5.1.0 or later switches do not indicate that a default configuration is enabled when you use the **cfgShow** or **cfgActvShow** commands. For more information about default zoning, "Configuring LSANs and Zoning" on page 11-22.

The following example procedure illustrates how LSANs control which devices can communicate with each another. The example procedure shows the creation of two LSANs (called *lsan_zone_fabric75* and *lsan_zone_fabric2*), which involve the following devices:

- Switch1 and the host are in fabric75
- Switch2, Target A, and Target B are in fabric2
- Switch1 is connected to the SilkWorm 48000 with an FR4-18i using an EX_Port or VEX_Port
- Switch2 is connected to the SilkWorm 48000 with an FR4-18i using another EX_Port or VEX_Port
- Host has WWN 10:00:00:00:c9:2b:c9:0c (connected to switch1)
- Target A has WWN 50:05:07:61:00:5b:62:ed (connected to switch2)
- Target B has WWN 50:05:07:61:00:49:20:b4 (connected to switch2)

The following procedure shows how to control device communication with LSAN.

To control device communication with LSAN

1. Log in as admin and connect to switch1.
2. Enter the **nsShow** command to list the WWN of the host (10:00:00:00:c9:2b:c9:0c). Note that the **nsShow** output displays both the port WWN and node WWN; the port WWN must be used for LSANs.

```
switch:admin_06> nsshow
{
 Type Pid       COS     PortName                NodeName                TTL(sec)
 N    060f00;   2,3;    10:00:00:00:c9:2b:c9:0c;   20:00:00:00:c9:2b:c9:0c; na
    FC4s: FCP
    NodeSymb: [35] "Emulex LP9002 FV3.91A3  DV5-5.20A6 "
    Fabric Port Name: 20:0f:00:05:1e:37:00:44
    Permanent Port Name: 10:00:00:00:c9:2b:c9:0c
The Local Name Server has 1 entry }
```

3. Enter the **zoneCreate** command to create the LSAN *lsan_zone_fabric75*, which includes the host.

```
switch:admin_06> zonecreate "lsan_zone_fabric75", "10:00:00:00:c9:2b:c9:0c"
```

4. Enter the **zoneAdd** command to Target A to the LSAN.

```
FID75Domain5:admin> zoneadd "lsan_zone_fabric75", "50:05:07:61:00:5b:62:ed"
```

5. Enter the **cfgAdd** and **cfgEnable** commands to add and enable the LSAN configuration:

```
switch:admin_06> cfgadd "zone_cfg", "lsan_zone_fabric75"
switch:admin_06> cfgenable "zone_cfg"
You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the
current configuration selected.
Do you want to enable 'zone_cfg' configuration  (yes, y, no, n): [no] y
zone config "zone_cfg" is in effect
Updating flash …
```

6. Log in as admin to fabric2.
7. Enter the **nsShow** command to list Target A (50:05:07:61:00:5b:62:ed) and Target B (50:05:07:61:00:49:20:b4).

```
switch:admin_06> nsshow
{
 Type Pid    COS     PortName                  NodeName                 TTL(sec)
 NL   0508e8; 3;     50:05:07:61:00:5b:62:ed;   50:05:07:61:00:1b:62:ed; na
    FC4s: FCP [IBM     DNEF-309170     F90F]
    Fabric Port Name: 20:08:00:05:1e:34:11:e5
    Permanent Port Name: 50:05:07:61:00:5b:62:ed
 NL   0508ef; 3;     50:05:07:61:00:49:20:b4;   50:05:07:61:00:09:20:b4; na
    FC4s: FCP [IBM     DNEF-309170     F90F]
    Fabric Port Name: 20:08:00:05:1e:34:11:e5
    Permanent Port Name: 50:05:07:61:00:49:20:b4
The Local Name Server has 2 entries }
```

8. Enter the **zoneCreate** command to create the LSAN *lsan_zone_fabric2*, which includes the host (10:00:00:00:c9:2b:6a:2c), Target A, and Target B.

```
switch:admin_06> zonecreate "lsan_zone_fabric2",
"10:00:00:00:c9:2b:c9:0c;50:05:07:61:00:5b:62:ed;50:05:07:61:00:49:20:b4"
```

9. Enter the **cfgShow** command to verify that the zones are correct.

```
switch:admin_06> cfgshow
Defined configuration:
 zone:  lsan_zone_fabric2
                10:00:00:00:c9:2b:c9:0c; 50:05:07:61:00:5b:62:ed;
                50:05:07:61:00:49:20:b4

Effective configuration:
 no configuration in effect
```

10. Enter the **cfgAdd** and **cfgEnable** commands to create and enable the LSAN configuration:

```
switch:admin_06> cfgadd "zone_cfg", "lsan_zone_fabric2"
switch:admin_06> cfgenable "zone_cfg"
You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the
current configuration selected.
Do you want to enable 'zone_cfg' configuration  (yes, y, no, n): [no] y
zone config "zone_cfg" is in effect
Updating flash ...
```

11. Log in as an admin and connect to the SilkWorm 48000 with a FR4-18i.
12. Enter the following commands to display information about the LSANs:
    - **lsanZoneShow -s** shows the LSAN.

```
switch:admin_06> lsanzoneshow -s
Fabric ID: 2 Zone Name: lsan_zone_fabric2
        10:00:00:00:c9:2b:c9:0c  Imported
        50:05:07:61:00:5b:62:ed  EXIST
        50:05:07:61:00:49:20:b4  EXIST
Fabric ID: 75 Zone Name: lsan_zone_fabric75
        10:00:00:00:c9:2b:c9:0c  EXIST
        50:05:07:61:00:5b:62:ed  Imported
```

    - **fcrPhyDevShow** shows the physical devices in the LSAN.

```
switch:admin_06> fcrphydevshow
        Device           WWN            Physical
        Exists                            PID
       in Fabric
       ----------------------------------------
           75 10:00:00:00:c9:2b:c9:0c c70000
           2  50:05:07:61:00:5b:62:ed 0100ef
           2  50:05:07:61:00:5b:62:ed 0100e8
 Total devices displayed: 3
```

    - **fcrProxyDevShow** shows the proxy devices in the LSAN.

```
switch:admin_06> fcrproxydevshow
        Proxy           WWN             Proxy      Device   Physical    State
        Created                           PID       Exists     PID
       in Fabric                                   in Fabric
      ------------------------------------------------------------------------
           2   50:05:07:65:05:84:09:0e  01f001       75     0100ef   Imported
           2   50:05:07:65:05:84:0b:83  01f000       75     0100e8   Imported
           75  10:00:00:00:c9:2b:c9:0c  02f000        2     c70000   Imported
Total devices displayed: 3
```

On the SilkWorm 48000 with an FR4-18i blade, the host and fabric75 are imported, because both are defined by *lsan_zone_fabric2* and *lsan_zone_fabric75*. However, target B defined by *lsan_zone_fabric75* is not imported because *lsan_zone_fabric2* does not allow it.

When a PLOGI, PDISC, or ADISC arrives at the SilkWorm 48000 with an FR4-18i blade, the SID and DID of the frame are checked. If they are LSAN-zoned at both SID and DID edge fabrics, the frame will be forwarded to the DID. If they are not zoned, only the PLOGI is dropped; for the remaining frames zoning enforcement takes place in the edge fabrics.

You can also set up LSAN zones using the following procedure.

To set up LSAN zones

1. Create a telnet or Web Tools connection to the SilkWorm switch.
2. Configure the LSAN, using the LSAN_xxxx naming schema, and append the new LSAN to the active zone set.
3. Enable zone configuration that now includes the new LSAN.

Configuring Backbone Fabrics for Interconnectivity

If you want devices in backbone fabrics to communicate with devices in edge fabrics, follow the steps in the preceding section, but instead of configuring the LSAN for the second edge fabric, do so for the backbone fabric.

Optional Configuration Procedures

This section describes optional configuration procedures; none are required to set up the FC-FC Routing Service.

FC Router Port Cost (Optional)

In Fabric OS v5.2.0, FCR optimizes the usage of the router port links by directing the traffic to the link with the smallest router port cost. The FC router port cost is similar to the link cost setting available on E_Ports, which allows the user to customize traffic flow. The router port link cost values are either 1000 or 10,000. The router module chooses the router port path based on the lowest cost per FID connection. If multiple paths exist with the same lowest cost, there will be load sharing over these paths.

You need only to differentiate between VEX and EX_Port link going to the same edge fabric as VEX and EX_Ports have very different speeds. For example, if there are VEX and EX_Port connections to the same edge fabric, the traffic is directed through the EX_Port link.

Every IFL has a default cost. The default router port cost values are:

- 1000 for legacy (v5.1 or XPath FCR) IFL
- 1000 for EX_Port IFL
- 10,000 for VEX_Port IFL

The FCR router port cost settings are 0, 1000 or 10,000. If the cost is set to 0, the default cost will be used for that IFL. The FC router port cost is persistent and is saved in the existing port configuration file.

Router port cost is passed to other routers in the same backbone. Link costs from the front domain to the translate (xlate) domain remain at 10,000. You can use the **lsDbShow** from the edge fabric to display these link costs.

Using Router Port Cost

Initially, the router port cost value for an EX_Port is set when the EX_Port is created when the **portCfgExport** command is issued. However, you can use the **fcrRouterPortCost** to modify the cost for that port. You can configure the EX or VEX_Port with either 1000 or 10,000 values. If you wish to differentiate between two EX_Port links with different speeds, you can assign 1000 to one link and 10,000 to another link.

For details about the use of any of the following commands, see the *Fabric OS Command Reference*.

To set and display the router port cost

1. Disable any port on which you want to set the router port cost.
2. Enable admin for the EX_Port/VEX_Port with **portCfgExport** or **portCfgVexport**.
3. Enter the **fcrRouterPortCost** command to display the router port cost per EX_Port.

```
switch:admin_06> fcrrouterportcost
 Port             Cost
------------------------
7/3              1000
7/4              1000
7/9              1000
7/10             1000
7/13             1000
10/0             1000
```

   You can also use the **fcrRouteShow** and **fcrDbgRoutShow** to display the router port cost.

4. Enter the **fcrRouterPortCost** command with a port and slot number, to display the router port cost for a single EX_Port.

```
switch:admin_06> fcrrouterportcost 7/10
Port             Cost
------------------------
7/10             1000
```

5. Enter the **fcrRouterPortCost** command with a port and slot number and a specific cost, to set the router port cost for a single EX_Port.

```
switch:admin_06> fcrrouterportcost 7/10 10,000
```

6. Enter the **fcrRouterPortCost** command**,** to set the cost of the EX_Port back to the default.

```
switch:admin_06> fcrrouterportcost 7/10 0
```

Upgrade, Downgrade, and HA Considerations

- For HA, the router port cost is synchronized to the standby CP.
- When upgrading the router switch from a prior Fabric OS version to Fabric OS v5.2.0, during warm recovery, if EX_Port is online, the default port cost is used.
- When downgrading the router switch from Fabric OS v5.2.0 to a prior Fabric OS version that does not support router port cost, the port configuration file retains the router port cost values. However, they are not used by the legacy Fabric OS.
- Legacy switches in the backbone fabric program all the router ports without considering router port cost. Fabric OS v5.2.0 considers the legacy router port cost as 1000 for both EX or VEX_ports.

Port Cost Considerations

The router port cost has the following considerations:

- Router port sets are defined as follows:
  - 0-7 and FCIP Tunnel 16-23
  - 8-15 and FCIP Tunnel 24-31

  More than two routers sets can exist on a chassis (e.g. SilkWorm 48000 Director).
- The router port cost does not help distinguish one IFL (or EX/VEX_Port link) from another, if all the IFLs are connected to the same port set. Therefore, if you connect IF1 and IFL2 to the same edge fabric in port set 0–7 and then configure them to different router port costs, traffic is still balanced across all the IFLs in the same port set.
- Use proper SAN design guideline to connect the IFLs to different port sets for effective router port cost use. For example, if both a low speed VEX_Port and a high-speed EX_Port are going to the same edge fabric, connect the lower router cost IFLs to one port group (for example ports 0–7) and the higher router cost IFLs to another port group (for example ports 8–15). For VEX_Ports, you would use ports in the range of 16-23 or 24-31.

You can connect multiple EX_Ports or VEX_Ports to the same edge fabric. The EX_Ports can all be on the same SilkWorm 7500 or 48000 with an FR4-18i, or they can be on multiple routers. Multiple EX_Ports create multiple paths for frame routing. Multiple paths can be used in two different, but compatible, ways:

- Failing over from one path to another
- Using multiple paths in parallel to increase effective data transmission rates

EX_Ports and VEX_Ports, when connected, are assigned different router port costs and traffic will flow only through the EX_Ports. Routing failover is automatic, but it can result in frames arriving out of order when frames take different routes. The SilkWorm 7500 and 48000 with an FR4-18i can force in-order delivery, although frame delivery is delayed immediately after the path failover.

Source EX_Ports can balance loads across multiple destination EX_Ports attached to the same edge fabric using exchange IDs from the routed frames as keys to distribute the traffic.

Setting a Proxy PID

When a SilkWorm 7500 or SilkWorm 48000 with an FR4-18i is first configured, the PIDs for the proxy devices are automatically assigned. Proxy PIDs (as well as phantom domain IDs) persist across reboots.

The most common situation in which you would set a proxy PID is when you replace a switch. If you replace the switch and want to continue using the old PID assignments, you can configure it to do so; this value remains in the system even if the blade is replaced. So you can minimize disruption to the edge fabrics by setting the proxy PIDs to the same values used with the old hardware.

The **fcrProxyConfig** command displays or sets the persistent configuration of proxy devices. Used with the **-s** *slot* option, it can also influence the assignment of the xlate domain port number (which is used to determine the Area_ID field of the PID) and the Port_ID field. Like the PIDs in a fabric, a proxy PID must be unique. If the *slot* argument results in a duplicate PID, it will be ignored. Proxy PIDs are automatically assigned to devices imported into a fabric, starting at f001. For Proxy IDs projected to a McDATA edge fabric in McDATA fabric mode, use valid ALPAs (lower 8 bits).

See the **fcrProxyConfig** command in the *Fabric OS Command Reference* for more details.

Use the **fcrXlateConfig** command to display or assign a preferred domain ID to a translate domain. See the **fcrXlateConfig** command in the *Fabric OS Command Reference* for more details.

Matching Fabric Parameters

By default, EX_Ports and VEX_Ports detect, autonegotiate, and configure the fabric parameters without user intervention.

You can optionally configure these parameters manually. Execute the **configure** command on a switch in the edge fabric to change the fabric parameters of a switch in the edge fabric. Execute the **portCfgEXPort** command to change the fabric parameters of an EX_Port on the SilkWorm 7500 or 48000 with a FR4-18i blade. If you want to change the fabric parameters of a VEX_Port, then use the **portCfgVEXPort** command.

If using the **-p** option with the **portCfgEXPort** command, note that the PID mode for the SilkWorm 7500 or SilkWorm 48000 with an FR4-18i (the backbone fabric PID mode) and the edge fabric PID mode do not need to match, but the PID mode for the EX_port or VEX_Port and the edge fabric to which it is attached must (the **-t** option disables the Negotiate fabric parameter feature). Otherwise, the PID mode is autonegotiated. The various edge fabrics may have different PID modes.

Fabric parameter settings must be the same on EX_Ports or VEX_Ports and on the fabrics to which they are connected: E_D_TOV (error-detect timeout value), R_A_TOV (resource-allocation timeout value), and PID format. You can set the PID format on an EX_Port when you configure an interfabric link.

The default values for E_D_TOV and R_A_TOV for an EX_Port or VEX_Port match those values on other SilkWorm switches. Only if you have adjusted these parameters for the fabric do you need to adjust them for an EX_Port or VEX_Port.

The default values for R_A_TOV and E_D_TOV are the recommended values for all but very large fabrics (ones requiring four or more hops) or high-latency fabrics (such as ones using long-distance FCIP links).

EX_Port Frame Trunking (Optional)

In Fabric OS v5.2.0, you can configure EX_Ports to use frame based trunking just as you do regular E_Ports. EX_Port frame trunking support is designed to provide the best utilization and balance of frames transmitted on each link between the FCR and the edge fabric. You should trunk all ports connected to the same edge fabrics.

Note

FCR front domain has a higher node WWN—derived from the FC router—than that of the edge fabric. Therefore, the FCR front domain initiates the trunking protocol on the EX_Port.

After initiation, the first port from the trunk group that comes online is designated as the master port. The other ports that come online on the trunk group are considered the slave ports. Adding or removing a slave port does not cause frame drop.

The restrictions for EX_Port frame trunking are the same as for E_Ports–all the ports must be adjacent to each other using the clearly marked groups on the front of the product.

EX_Port frame-based trunking has a master trunk link. If the master trunk link goes down, the entire EX_Port-based trunk will reform and it will be taken offline for a short period of time. If there are no other links to the edge fabric from the backbone, the master port going offline may cause a traffic disruption in the backbone.

If router port cost is used with EX_Port trunking the master port and slave ports share the router port of the master port.

For information about setting up E_Port trunking on an edge fabric, see Chapter 19, “Administering ISL Trunking” in this guide.

Supported Configurations and Platforms

The EX_Port trunking is an FCR software feature and requires that you have the FCR trunking license installed on the FCR switch and on the edge fabric connected to the other side of the trunked EX_Ports. EX_Port trunking is supported only with Brocade edge fabrics. You can use EX_Port frame trunking in the following configurations and cases:

- With ports with speeds of 2 Gbit/sec up to a maximum speed of 4 Gbit/sec and trunking over long distance.
- In the edge fabric, when FCR is connected to a Condor-based switch (for example, the SilkWorm 4100, 4900, or FC 4-48 blade) that supports eight ports from the trunkable group.
- When FCR is connected to Bloom-based switches (for example, the SilkWorm 3800, 3900, 24000, or FC-16 blade) supporting four ports in the trunk group. If the edge fabric is a Bloom-based switch, the FCR will setup the trunking for four ports per trunk.
- When FCR is connected to an edge fabric using a mix of trunked and non-trunked EX_Ports. All will share the same front domain.
- With the Secure Fabric OS edge fabric.

High Availability Support

The EX_Port frame trunking feature also is a High Availability (HA) supported feature. The HA protocol for EX_Port trunking is as follows:

- If trunking is disabled prior to the HA failover, it remains disabled after the HA failover.
- If trunking is enabled prior to the HA failover, it remains enabled after the HA failover.

Backward Compatibility Support

For backward compatibility, FCR that supports EX_Port trunking can continue to interoperate with older FCR switches and all previously supported Brocade switches in the Brocade edge fabric.

Upgrade and Downgrade Considerations

The following table describes the upgrade and downgrade considerations for EX_Port Frame Trunking.

Table 11-1 Trunking Upgrade and Downgrade Considerations

Upgrade or Downgrade	Consideration
A firmware downgrade from Fabric OS v5.2.0 to Fabric OS v5.1.0	If EX_Port trunking is on, prior to the firmware downgrade, the script displays a message requesting that you disable EX_Port trunking.
A firmware upgrade from Fabric OS v5.1 to Fabric OS v5.2.0	The trunking feature is turned off for each EX_Port configured on the FCR switch. To explicitly enable trunking on all/one EX_Port(s), you can use the **portCfgTrunkPort** and **switchCfgTrunk** commands. See the following sections for examples.

Using EX_Port Frame Trunking

With EX_Port frame trunking, you can use the same CLI commands as you do for E_Port trunking.

Security Considerations

Administration control for EX_Port trunking is available through root, admin, and switch admin access.

Trunking Commands

The procedure for administering EX_Port frame trunking is the same as for E_Port trunking. You will need to initialize trunking on ports with **portCfgTrunkPort** or **switchCfgtrunk**, monitor traffic with the **portPerfShow** command. You need a Brocade license to enable trunking, however, you can disable trunking on the ports without the trunking license installed. You can use the following commands to enable and disable trunking:

- Disable (0) or enable (1) trunking on a single port using the **portCfgTrunkPort** command. The following example shows how to disable trunking on a single port and then how to enable trunking:.

```
switch:admin_06> portcfgtrunkport 1/10 0

switch:admin_06> portcfgtrunkport 1/10 1
```

- Disable (0) or enable (1) trunking on all the ports of a switch using the **switchCfgTrunk** command.The following example shows how to disable trunking on all ports and then how to enable trunking:

```
switch:admin_06> switchcfgtrunk 0

switch:admin_06> switchcfgtrunk 1
Configuration applied to all ports except the following VE/VEX_Ports (ports 128 -
143, 192 - 207).
```

When either command is executed to update the trunking configuration, the ports for which the configuration applies are disabled and reenabled with the new trunk configuration. As a result, the traffic through these ports might be disrupted for a short period of time. In addition to the commands for enabling and disabling trunking, you can also use the following E_Port commands for administering EX_Port Frame Trunking:

- Use **portCfgSpeed** and **switchCfgSpeed** to set speed for a port or switch
- Display lists of trunks and members of trunks with the **trunkShow** command
- Use **trunkDebug** to list link characteristics for troubleshooting.
- Use the **switchShow** command to display the trunking information on the EX_Ports.

For details about using these commands, see Chapter 19, “Administering ISL Trunking”and individual commands in the *Fabric OS Command Reference*.

To display EX_Port trunking information

1. Log in as an admin and connect to the switch.
2. Enter the **switchShow** command to display trunking information for the EX_Ports.

```
fcr_switch:admin_06> switchshow
```

The following is an example of a master EX_Port and a slave EX_Port displayed in **switchShow**.

```
Index Slot Port Address Media Speed State
==========================================
16    2    0    ee1000   id    N4   No_Light
17    2    1    ee1100   id    N4   Online    EX_Port  (Trunk port, master is
Slot  2 Port  2 )
18    2    2    ee1200   id    N4   Online    EX_Port  10:00:00:05:1e:35:bb:32
"MtOlympus_82" (fabric id = 2 )(Trunk master)
19    2    3    ee1300   id    N4   No_Light
20    2    4    ee1400   id    N4   Online    EX_Port  (Trunk port, master is
Slot  2 Port  7 )
21    2    5    ee1500   id    N4   Online    EX_Port  (Trunk port, master is
Slot  2 Port  7 )
22    2    6    ee1600   id    N4   Online    EX_Port  (Trunk port, master is
Slot  2 Port  7 )
23    2    7    ee1700   id    N4   Online    EX_Port  10:00:00:60:69:80:1d:bc
"MtOlympus_72" (fabric id = 2 )(Trunk master)
```

To display trunking information for a switch

1. Use the **trunkShow** command to display trunking information for a switch.

```
switch:admin> trunkshow
 1: 6 ->  4   10:00:00:60:69:51:43:04   deskew 15   MASTER

 2:15 -> 13   10:00:00:60:69:51:43:04   deskew 16   MASTER
   12 -> 12   10:00:00:60:69:51:43:04   deskew 15
   14 -> 14   10:00:00:60:69:51:43:04   deskew 17
   13 -> 15   10:00:00:60:69:51:43:04   deskew 16

 3:24 -> 14   10:00:00:60:69:51:42:dd   deskew 15   MASTER
```

Monitoring Resources

It is possible to exhaust resources, such as proxy PIDs. Whenever a resource is exhausted, Fabric OS generates an error message. The messages are described in the *Fabric OS Message Reference*.

You can monitor SilkWorm 7500 and 48000 resources using the **fcrResourceShow** command. The **fcrResourceShow** command shows FCR resource limits and usage and includes:

- LSAN Zones and LSAN devices—The information shows the maximum versus the currently used zones and device database entry. LSAN is the count of LSANs in all the edge fabrics. If they are defined in two edge fabrics, they are counted as two and not one. One device imported into multiple edge fabrics counts multiple times.
- Proxy Device Slots—The display shows the maximum versus the currently used proxy device slots. A proxy device is presented to an edge fabric as being connected to a translate domain *slot*. A slot is the port number and AL_PA combination. The slot-to-device WWN association is persistently stored. The physical and proxy devices use the 10000 device slots.
- Displays the maximum pool size for translate phantom node and port WWNs and shows the number of translate node and port WWNs from this pool.

- Phantom Node WWN—The display shows the maximum versus the currently allocated phantom switch node WWNs. The phantom switch requires node WWNs for SFPF and manageability purposes. Phantom node names are allocated from the pool sequentially and are not reused until the pool is exhausted and rolls over. The last allocated phantom node WWN is persistently stored.

  If the switch is disabled, the phantom node WWNs are not returned to the pool until the system reboots, because the phantom switch could still be accessible through other switches. Across a switch reboot, the allocation starts from the next usable WWN from the pool and not from the beginning.

- Phantom Port WWNs—The display shows the maximum versus the currently used phantom domain port WWNs. Phantom domain ports require port WWNs for manageability purposes. Phantom domain ports include ports connecting front and translate domains (virtual ISLs), translate domain ports for proxy devices, and EX_Ports. Phantom port names are allocated from the pool sequentially and are not resumed until the pool is exhausted and rolls over. The last allocated phantom port WWN is persistently stored.

  If the switch is disabled, phantom port WWNs are not returned to the pool until the system reboots, because the phantom switch might still be accessible through other switches. Across the switch reboot, the allocation starts from the next usable WWN base from the pool and not from the beginning.

- Proxy Devices—The maximum versus the currently used proxy devices.
- NR_Port—The maximum versus the currently used NR_Port entries. Destination NR_Port entries are stored at every physical port for routing decision purposes.

The following example shows the use of the **fcrResourceShow** command display per physical port (EX_Port) resources, which includes:

```
switch:admin> fcrresourceshow
Daemon Limits:
Max Allowed Currently Used
---------------------------------
LSAN Zones: 3000 28
LSAN Devices: 10000 51
Proxy Device Slots: 10000 20
WWN Pool Size Allocated
---------------------------------
Phantom Node WWN: 8192 5413
Phantom Port WWN: 32768 16121
Port Limits:
Max proxy devices: 2000
Max NR_Ports: 1000
Currently Used(column 1: proxy, column 2: NR_Ports):
0 | 0 34
1 | 3 34
4 | 0 0
5 | 0 0
6 | 0 0
7 | 0 0
8 | 6 34
9 | 6 34
10 | 6 34
11 | 6 34
12 | 6 34
13 | 6 34
14 | 6 34
15 | 6 34
16 | 8 34
17 | 8 34
18 | 8 34
19 | 8 34
20 | 8 34
21 | 8 34
22 | 8 34
23 | 8 34
```

See the *Fabric OS Command Reference* for details about the **fcrResourceShow** command.

Routing ECHO

The FC-FC Routing Service enables you to route the ECHO generated when an **fcPing** command is issued on a switch, providing **fcPing** capability between two devices in different fabrics across the SilkWorm 7500 or 48000 with an FR4-18i blade.

To check for Fibre Channel connectivity problems

1. On the edge Fabric OS switch, make sure that the source and destination devices are properly configured in the LSAN zone before entering the **fcPing** command, which:
   - Checks the zoning configuration for the two ports specified.
   - Generates an ELS (Extended Link Service frame) ECHO request to the source port specified and validates the response.
   - Generates an ELS ECHO request to the destination port specified and validates the response.

```
switch:admin_06> fcping 0x060f00 0x05f001
Source:         0x60f00
Destination:    0x5f001
Zone Check:     Zoned

Pinging 0x60f00 with 12 bytes of data:
received reply from 0x60f00: 12 bytes time:501 usec
received reply from 0x60f00: 12 bytes time:437 usec
received reply from 0x60f00: 12 bytes time:506 usec
received reply from 0x60f00: 12 bytes time:430 usec
received reply from 0x60f00: 12 bytes time:462 usec
5 frames sent, 5 frames received, 0 frames rejected, 0 frames timeout
Round-trip min/avg/max = 430/467/506 usec

Pinging 0x5f001 with 12 bytes of data:
received reply from 0x5f001: 12 bytes time:2803 usec
received reply from 0x5f001: 12 bytes time:2701 usec
received reply from 0x5f001: 12 bytes time:3193 usec
received reply from 0x5f001: 12 bytes time:2738 usec
received reply from 0x5f001: 12 bytes time:2746 usec
5 frames sent, 5 frames received, 0 frames rejected, 0 frames timeout
Round-trip min/avg/max = 2701/2836/3193 usec
```

2. Regardless of the device's zoning configuration, the **fcPing** command sends the ELS frame to the destination port. A destination device can take any one of the following actions:
   - Send an ELS Accept to the ELS request.
   - Send an ELS Reject to the ELS request.
   - Ignore the ELS request.

   There are some devices that do not support the ELS ECHO request. In these cases, the device will either not respond to the request or send an ELS reject. When a device does not respond to the ELS request, further debugging is required; however, do not assume that the device is not connected.

   For details about the **fcPing** command, see the *Fabric OS Command Reference*.

Upgrade and Downgrade Considerations

If you downgrade to a version of Fabric OS that does not support FC-FC Routing Services, then your FC-FC routing configuration will be lost. Brocade recommends that you enter the **configUpload** command to save your FC-FC routing configuration before performing any downgrades.

If you have a SilkWorm 48000 with an FR4-18i blade configured using chassis option 5 with the blade powered on and you attempt to downgrade to an earlier version of the Fabric OS, the downgrade will not occur until *after* you power off and remove the FR4-18i blade. If you attempt to swap out the FR4-18i blade with another (for example, the FC4-16 or FC4-32), the port's configuration will be cleared; however, all other port configurations are retained.

Interoperability with Legacy FCR Switches

The following interoperability considerations apply when administering legacy FCR switches in the same backbone (BB) fabric as switches supporting Fabric OS v5.2.0:

- When a legacy switch is connected to the fabric, a RAS log message is issued indicating that the capability of the backbone (BB) fabric is lower as legacy FCR switches (those with XPath OS and Fabric OS v5.1) support lower capability limits.
- If an XPath OS switch joins the BB fabric, a RAS log message is issued indicating that the backbone-to-edge (B2E) fabric is not supported with the XPath OS switch.

For the exact RAS log message descriptions, see the following RAS logs: FCR_1055, FCR_1056, and FCR_1073.

Backward Compatibility

In a fabric with Secure Fabric OS enabled, the edge fabric must have Fabric OS v3.2, v4.4.0, or later because only DH-CHAP authentication is supported.

For a nonsecure fabric, the hardware and firmware compatibility is described in Table 11-2.

Table 11-2 Hardware and Firmware Compatibility for Nonsecure Fabrics

Model	Supported Version
SilkWorm 2000 series	Fabric OS v2.6.1 or later
SilkWorm 3200, 3600, 3800	Fabric OS v3.1.0 or later
SilkWorm 3900, 12000	Fabric OS v4.1.0x or later
SilkWorm 3250, 3850, 24000	Fabric OS v4.2.0x or later
SilkWorm 4100	Fabric OS v4.4.0 or later
SilkWorm 200E, 48000	Fabric OS v5.0.1 or later
SilkWorm 4900, 7500	Fabric OS v5.1.0 or later
SilkWorm 48000 with FC4-48 or FC4-16IP blades	Fabric OS v5.2.0 or later

Connecting to McDATA SANs

Fabric OS v5.2.0 lets you connect a Brocade fabric to a McDATA fabric. Because of the high degree of connectivity Brocade provides, the devices across the remote fabrics can be shared. Fabric OS v5.2.0 furnishes the FC router with the ability to connect to McDATA fabrics in both McDATA Open mode or McDATA Fabric mode.

Note

Open mode is for use only by McDATA switches. Brocade switches in open mode can not be part of the McDATA edge fabric.

Connected SANs provide additional functionality not possible with segregated SANs. Some of these include:

- Island consolidation—Use the Fabric OS v5.2.0 FC router to connect isolated McDATA and Brocade Fabrics and share devices.
- Backup consolidation—Consolidate backup solutions across Brocade and McDATA fabrics.
- Manageable large scale storage network—Use the Fabric OS v5.2.0 FC router to localize traffic while connecting devices in the metaSAN. This provides a large number of fabrics with a large number of devices.
- Sharing across a FCIP link - share devices between Brocade and McDATA fabric over campus Ethernet or over long-distance IP link beyond 1000 km.
- Sharing across a long-distance FC link - share devices between Brocade fabrics over long-distance FC links as far as 300 km.
- LUN sharing - use your high-end RAID array connected to a McDATA director to share targets with a Brocade fabric; just connect one McDATA director port to an FC router EX_Port and the one EX_Port to the Brocade fabric.

The connectivity limitations of a metaSAN comprising Brocade and McDATA fabrics are defined by the scalability of each individual fabric. The latest scalability information can be found at the Brocade Connect Web site at *www.brocade.com*. Refer to the McDATA fabric documentation for McDATA scalability considerations.

Supported Platforms and Configurations

Connection to McDATA SANs through an FC router is supported for the following Brocade hardware platforms:

- SilkWorm 7500 switch
- SilkWorm 48000 director with an FR4-18i blade

The following McDATA configurations are supported through the Fabric OS Fibre Channel Routing Service:

- Interoperability with McDATA E/0S 7.x and 8.x
- McDATA 3x16, 3x32, 4300, 4500, 6064 (1 Gbit and 2 Gbit), 6140 switch and director platforms
- Up to a 400-port fabric
- Up to 8 domains

Check for the latest list of tested and approved equipment at the Brocade Connect Web site at *www.brocade.com*.

Configuring the Fabrics for Interconnectivity

When connecting a Brocade fabric with a McDATA fabric using an FC Router, you must configure the switch on both fabrics as well as router described in:

- "Connectivity Modes," next
- "Configuring the FC Router" on page 11-40
- "Configuring McDATA for Interconnection" on page 11-44
- "Completing the Configuration" on page 11-48

Note
Trunking is not supported on EX_Ports connected to the McDATA fabric.

Connectivity Modes

You can connect to McDATA fabrics in both McDATA Open mode or McDATA Fabric mode. If the mode is not configured correctly, the port is disabled for incompatibility.

To allow interconnectivity with McDATA SANs, the command line interface (CLI) command **portCfgExPort** uses the **-m** option to indicate the connectivity mode. Table 11-3 lists the valid parameters to use with the **-m** option to set the connectivity mode.

Table 11-3 portCfgExPort -m Values

Value	Description	Use
0	Brocade Native	Default mode
1	McDATA Open mode	When the neighboring McDATA switch is running in open mode
2	McDATA Fabric mode (native mode)	When the neighboring McDATA switch is running in native mode
3	McData Fabric legacy mode	Not currently used.

You can display the current operational mode of the EX_Port by issuing the **portCfgExPort** command with the port number as the only parameter.

See "Configuring an Interfabric Link" on page 11-16 for details about the **portCfgExPort** command, which is used for McDATA Fabric mode on Fabric OS v5.2.0.

Configuring the FC Router

When configuring a Brocade fabric to connect to a McDATA fabric, you must perform some preparation on the FC router. The following procedure shows how to connect an EX_Port of an FC router fabric to a McDATA fabric configured in Fabric mode.

1. Verify the McDATA firmware version. To display the front domain on the McDATA fabric, use the McDATA **show** command.
2. Using the **version** command, make sure that Fabric OS v5.2.0 is installed on the SilkWorm 7500 or FR4-18i blade, as shown in the following.

```
fcr_switch:AD255:admin> version
Kernel:     2.4.19
Fabric OS:  v5.2.0
Made on:    Thu Sep 21 01:15:34 2006
Flash:      Fri Sep 22 20:53:48 2006
BootProm:   4.5.3
```

3. On the Silkworm 7500 and FR4-18i blade, disable the port being configured as an EX_Port (the one connected to the SilkWorm) by issuing the **portDisable** command.

   You can verify that the port has been disabled by issuing the **portShow** command for the port.
4. Using the **portCfgExPort** command, configure the port as an EX_Port and provide a Fabric ID (FID).

   If no FID is specified, multiple links to the same fabric are assigned different FIDs. This results in a fabric ID oversubscription.

   For the SilkWorm fabric, use the **-p** flag of the **portCfgExPort** command to match the PID setting for the fabric to which you are connecting. This port will connect to the SilkWorm switch.

   The following example sets port 10/12 to admin-enabled, assigns a Fabric ID of 41 and sets the port to Core PID and to Brocade mode. For complete information about any Fabric OS command, see Appendix B, “Configuring Interoperability Mode”.

```
switch:admin_06> portcfgexport 10/12 -a 1 -f 41 -p 1 -m 0
```

5. Re-enable the port by issuing the **portEnable** command.

```
Switch:admin_06> portenable 10/12
```

6. Still on the Silkworm 7500 and FR4-18i blade, use the **portDisable** to disable the EX_Port that you will use to connect to the McDATA switch Ports are persistently disabled by default.

```
Switch:admin_06> portdisable 10/13

Switch:admin_06> switchshow
switchName:     b48000_5x_1/PORT/P
switchType:       42.2
switchState:      Online
switchMode:     Native
switchRole:       Principal
switchDomain:  3
switchId:           fffc03
switchWwn:       10:00:00:60:69:e4:00:86
zoning:             ON (test)
switchBeacon:   OFF
blade3 Beacon:  OFF
blade8 Beacon:  OFF
blade10 Beacon: OFF
FC Router:        ON
FC Router BB Fabric ID:            1

Index Slot Port Address Media Speed State     Proto

==================================================

112   10    0   037000   id    N4   No_Light         Disabled (Persistent)
113   10    1   037100   id    N4   No_Light         Disabled (Persistent)
114   10    2   037200   id    N4   No_Light         Disabled (Persistent)
115   10    3   037300   id    N4   No_Light         Disabled (Persistent)
116   10    4   037400   id    N4   No_Light         Disabled (Persistent)
117   10    5   037500   id    N4   No_Light         Disabled (Persistent)
118   10    6   037600   id    N4   No_Light         Disabled (Persistent)
119   10    7   037700   id    N4   No_Light         Disabled (Persistent)
120   10    8   037800   id    N4   No_Sync          Disabled (Persistent)
121   10    9   037900   id    N2   No_Light
122   10   10   037a00   id    N2   No_Light
123   10   11   037b00   id    N2   No_Light
124   10   12   037c00   id    N2   Online           EX-Port  10:00:00:05:1e:34:e0:6a
                                                   "b4100_5x_1"(fabric id = 41 )
125   10   13   037d00   id    N2   In_Sync          Disabled
126   10   14   037e00   id    N2   No_Light
127   10   15   037f00   id    N2   No_Light
240   10   16   03f000   --    --   Offline          Disabled (Persistent)
241   10   17   03f100   --    --   Offline          Disabled (Persistent)
242   10   18   03f200   --    --   Offline          Disabled (Persistent)
243   10   19   03f300   --    --   Offline          Disabled (Persistent)

...this example has been truncated
```

7. Issue the **portCfgExPort** command to configure the port as an EX_Port with a different FID within the McDATA fabric PID mode.

   This port can now connect to a McDATA switch in McDATA Fabric mode or McDATA Open mode.

   The following example sets port 10/13 to admin-enabled, assigns a Fabric ID of 37, and sets the McDATA connection mode to McDATA fabric.

```
switch:admin_06> portcfgexport 10/13 -a 1 -f 37 -m 2
```

8. Enable the port by issuing the **portEnable** command.

```
switch>:admin_06 portenable 10/13
```

   - Connect IFL 1 and verify EX_PORT connectivity. Repeat for all Brocade fabric IFLs.
   - Connect IFL (n) for the McDATA fabric and verify EX_PORT connectivity. Repeat for all McDATA fabric IFLs.

9. Log in to the FC router and issue the **switchShow** command to display the McDATA switch that is connected to the FC router EX_Port. The following example illustrates the use of the **switchShow** command.

   The output from the switchshow command is show in the following example:

```
switch>:admin_06 > switchshow
switchName:     b48000_5x_1
switchType:     42.2
switchState:    Online
switchMode:     Native
switchRole:     Principal
switchDomain:   3
switchId:       fffc03
switchWwn:      10:00:00:60:69:e4:00:86
zoning:         ON (test)
switchBeacon:   OFF
blade3 Beacon:  OFF
blade8 Beacon:  OFF
blade10 Beacon: OFF
FC Router:      ON
FC Router BB Fabric ID:            1

Index Slot Port Address Media Speed State     Proto
==================================================
112   10    0   037000   id    N4   No_Light         Disabled (Persistent)
113   10    1   037100   id    N4   No_Light         Disabled (Persistent)
114   10    2   037200   id    N4   No_Light         Disabled (Persistent)
115   10    3   037300   id    N4   No_Light         Disabled (Persistent)
116   10    4   037400   id    N4   No_Light         Disabled (Persistent)
117   10    5   037500   id    N4   No_Light         Disabled (Persistent)
118   10    6   037600   id    N4   No_Light         Disabled (Persistent)
119   10    7   037700   id    N4   No_Light         Disabled (Persistent)
120   10    8   037800   id    N4   No_Sync          Disabled (Persistent)
121   10    9   037900   id    N2   No_Light
122   10   10   037a00   id    N2   No_Light
123   10   11   037b00   id    N2   No_Light
124   10   12   037c00   id    N2   Online          EX-Port  10:00:00:05:1e:34:e0:6a "b4100_5x_1"
                                                   (fabric id = 41 )
125   10   13   037d00   id    N2   Online          EX-Port  10:00:08:00:88:03:53:fc "mcd_37"
                                                   (fabric id = 37 )
126   10   14   037e00   id    N2   No_Light
127   10   15   037f00   id    N2   No_Light
240   10   16   03f000   --    --   Offline          Disabled (Persistent)
241   10   17   03f100   --    --   Offline          Disabled (Persistent)
242   10   18   03f200   --    --   Offline          Disabled (Persistent)
243   10   19   03f300   --    --   Offline          Disabled (Persistent)
...this example has been truncated
```

10. You can now physically attach your ISLs from the Silkworm 7500 and FR4-18i blade to the switches.

    ISLs apply only to Brocade switches that are not connected as an edge fabric (IFLs). When a McDATA switch is present, it is assumed that you are creating an edge fabric.

    For information about Brocade edge fabric setup on E_ports and interswitch linking, see Chapter 19, "Administering ISL Trunking". For information on EX_Port Frame trunking setup on the FCR switch, see "Using EX_Port Frame Trunking" on page 11-31.

11. Capture a SAN profile of the McDATA and Brocade SANs, identifying the number of devices in each SAN.

    By projecting the total number of devices and switches expected in each fabric when the LSANs are active, you can quickly determine the status of the SAN by issuing the commands **nsAllShow** and **fabricShow** on the Brocade fabric. The **nsAllShow** displays the global name server information and **fabricShow** displays the fabric membership information. The following examples illustrate the use of these commands.

```
wta_switch:admin> fabricshow
Switch ID   Worldwide Name           Enet IP Addr    FC IP Addr      Name
-------------------------------------------------------------------------
  1: fffc01 10:00:00:60:69:e4:00:36 10.33.13.91     220.220.220.91
"wtQA_saturn91"
218: fffcda 10:00:00:60:69:90:02:b3 10.33.13.218    10.10.10.11
>"SW3900_WT218"
```

    You can use SAN Pilot or EFCM to gather similar information for the McDATA fabric.

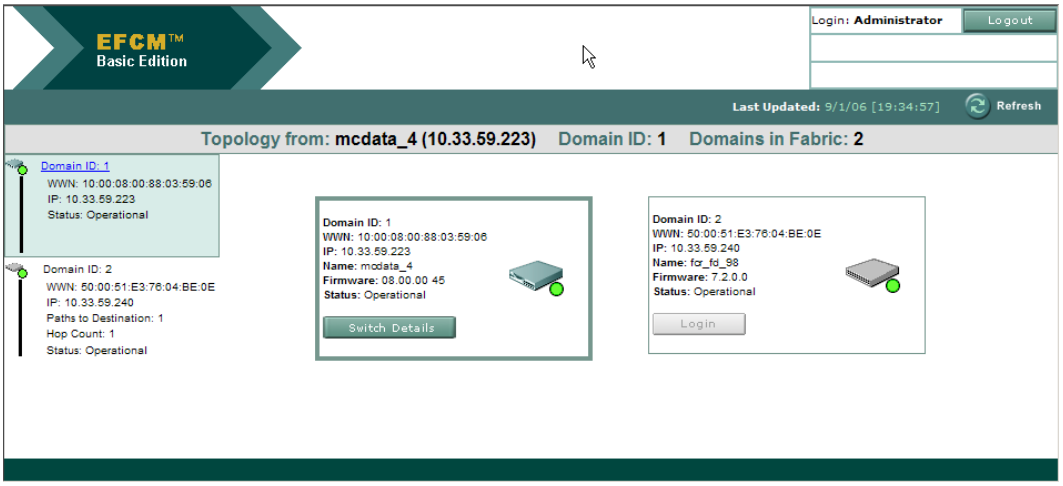

Figure 11-5 EFCM SAN Status

When you have configured the Silkworm 7500 and FR4-18i blade to connect to a Brocade fabric, you must create LSAN and zones for the SAN. Once you have set up LSAN zoning, you can issue the **cfgShow** command to verify that the zoning is correct.

Note

The procedures related to McDATA that are described in this section were current when the document was written, but may have changed since then. For the most up-to-date information, refer to the McDATA documentation about zone configuration on McDATA fabrics.

Configuring McDATA for Interconnection

To ensure connectivity with the Brocade fabric, you must prepare the McDATA fabric for use with the Silkworm 7500 and FR4-18i blade.

Note

The procedures described in this section were current when the document was written, but may have changed since then. For the most up-to-date information, refer to the McDATA documentation about zone configuration on McDATA fabrics.

To prepare the McDATA fabric

1. Log in to SAN Pilot or basic EFC Manager depending upon the firmware release.
2. From the SAN Pilot left navigation menu, select **Configure**.
3. Select the **Zoning** tab, then select the **Zones** tab. (select **Configure > Zoning** on EFCM).

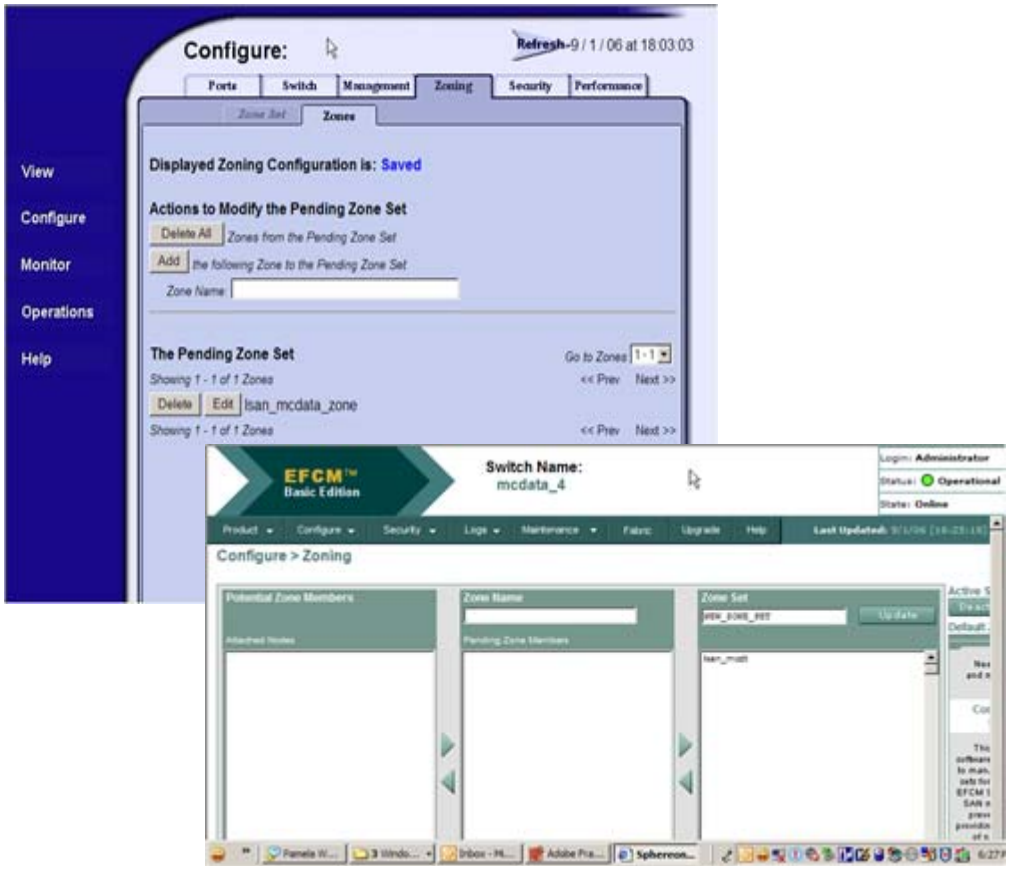

Figure 11-6 SAN Pilot and EFCM Zone Screens

Note

The screens provided in this section are for illustrative purposes only. Depending on the McDATA firmware release you are using, the McDATA web-based management tool may display a user interface different from those shown.

4. Type the desired name in the **Zone Name** field, using the LSAN_xxxx naming schema.

   In EFCM, move to the **Zone Name** field, and enter the desired name using the LSAN_xxxx naming schema.

5. In SAN Pilot, click the **Add** button to add the specified Zone. As shown in the following illustration, when you add the new zone name, the name is displayed in the **Pending Zone Set** list.

6. To add devices that are connected to the Brocade fabric, select **Edit** button in the **Pending Zone** set.

7. In the **Modify Zone** tab, enter the device WWN into the World Wide Name field and click the **Add** button. The **Pending Zone Membership List** is updated with the new Zone members.

   If you are using EFCM, in the WWN field of **Potential Zone Members** > **New Member**, enter the WWN port name and click **Add**.

Note

If you are using EFCM, 8.0 or later, some of the steps will be different. To obtain current information about your McDATA product, refer to http://www.mcdata.com and become familiar with the McDATA zoning user manual for your specific EFCM release.

8. Select the **Zone Set** tab in SAN Pilot.

   If you are using EFCM or the **Zoneset** Library window, tab to **Zone Sets** and select **File > New**.

9. Enter a name for the Pending Zone set in **Zone Set Name** and select **Rename**.

   The steps for EFCM will be similar.

10. Select the **Save** and **Activate the Pending Zone** set. The zones in the Pending Zone set are added to the Zone Set that you specified.

    In EFCM, return to the main window and select **Configure**, then select **Activate Zone Set** to launch the zone set activation window. Highlight the zone set to be activated and click **Next**. Click **Next** again, then **Start** to activate the zone set.

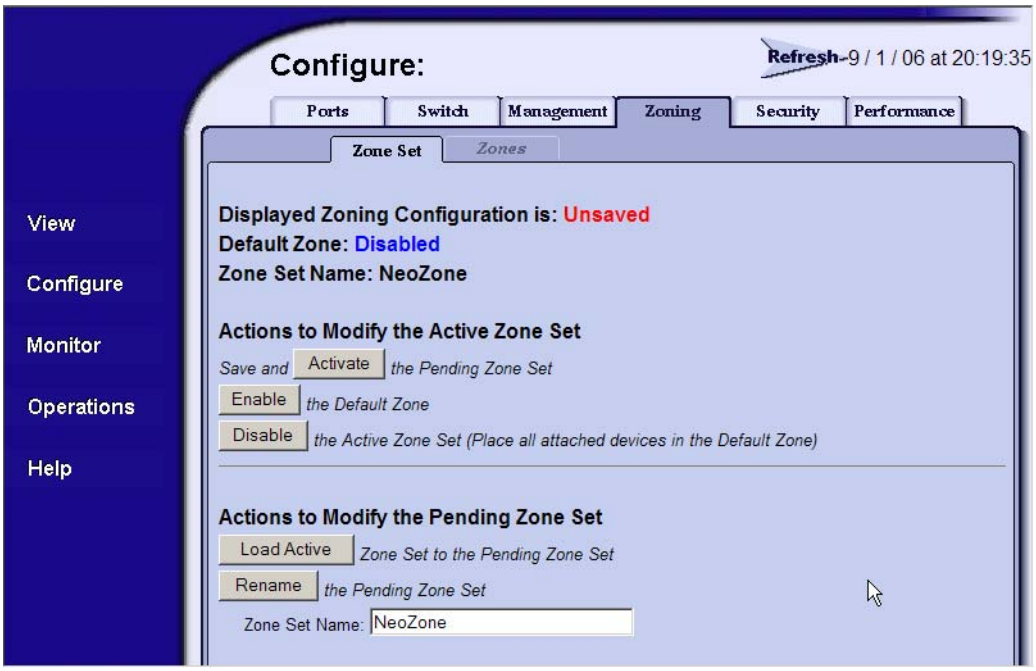

Figure 11-7 Adding a Zone Set Name in SAN Pilot

Regardless of the method used, you should now verify that the new zone set containing your LSAN has been added.

Alternately, use the following procedure:

1. Create the LSAN, using the LSAN_xxxx naming schema.
2. Append the newly-created zone set to a currently active zone set.
3. Activate the updated zone set.

LSAN Zoning with McDATA

An LSAN is defined by a zone in an edge fabric. When zoning an LSAN containing multiple manufacturers' fabrics (such as a Brocade-McDATA SAN), you must use port WWN. Because port IDs are not necessarily unique across fabrics, you cannot use the "domain,port" method of identification.

If the LSAN is configured and the proxy devices are created, the proxy device will show in the name server of the edge fabric and the xlate domain will show in the fabric of the edge fabric.

For more details about LSAN zoning, see "Configuring LSANs and Zoning" on page 11-22.

If the LSAN devices appear in only one of the fabrics in a multiple-fabric SAN, use the following procedure to correct the problem.

To correct errors if LSAN devices appear in only one of the fabrics

1. Log in to each fabric and verify that all of the devices are physically logged in.
2. Verify that the devices are properly configured in the LSAN zone in both edge fabrics.
3. Issue the **fabricShow** command on the Brocade fabric.
4. Use McDATA Enterprise Fabric Connectivity Manager (EFCM) or SAN Pilot to verify the McDATA fabric, including the front and translation domains.
5. Move back to the Silkworm 7500 and FR4-18i blade and issue the **fcrProxyDevShow** command on to verify that the devices are configured and exported.

```
switch:admin_06> fcrproxydevshow
  Proxy           WWN             Proxy      Device   Physical    State
 Created                           PID       Exists     PID
in Fabric                                   in Fabric
--------------------------------------------------------------------------------
    10   20:00:00:01:73:00:59:dd  05f001       12      610902   Imported
    10   21:00:00:e0:8b:04:80:76  02f002       11      340713   Imported
    10   50:06:01:68:40:04:d3:95  02f001       11      660713   Imported
    11   10:00:00:00:c9:2d:3d:5c  020001       10      011500   Imported
    11   50:06:01:60:40:04:d3:95  020002       10      011400   Imported
router:admin>
```

6. Create a telnet connection and configure the connection to capture text.
7. Issue the **supportShow** command, and save the output.
8. If the fabric does not appear:
   a. Disable the EX_Port on the connected fabric.
   b. Issue the **portLogClear** command for the port.
   c. Enable the port on the Silkworm 7500 and FR4-18i blade.
   d. Issue the **portLogDump** command for the port, capturing the output.

      Use the **portLogDump** tool to troubleshoot the problem, using the command output.

   If an EX_Port connecting a Silkworm 7500 and FR4-18i blade and an edge fabric is disabled due to an error, the error causing that port's most recent disabled state appears in the **switchShow** command output. This error appears until that port comes back online, even after the cables have been detached from the port.

   To remove the error listing in the **switchShow** output, reboot the Silkworm 7500 and FR4-18i blade. An example of the type of error displayed is ‘Incompatible port mode’.

Completing the Configuration

After you prepare the McDATA switch and the Silkworm 7500 and/or FR4-18i blade for use, complete the configuration using the following procedure.

1. Physically connect the EX_Port that you configured for the SilkWorm switch to the FC router.
2. Log in to the SilkWorm switch as an admin.
3. Physically connect the configured FC router EX_Port to the McDATA switch, and issue the **switchShow** command.

   New domains should be visible for each IFL (front domain) that connects the SilkWorm to the FC router, and one domain for the translate domain.
4. Start SAN Pilot (or EFCM) and select the fabric for the McDATA switch.
5. View the fabric topology.

   New domains should be visible for every FCR connection to the McDATA switch to the Silkworm 7500 and FR4-18i blade, and one domain for the translate domain.

   In EFCM, the McDATA switch should appear green, and the front domains (as well as the translate domain) are grayed out and inaccessible: EFCM cannot manage them. Tab to **Zone** and verify that the zone set configuration is correct: a blue icon beside each entry indicates that the devices are logged in to the fabric.

6. Log in to the SilkWorm switch and issue the **nsAllShow** or the **nsCamShow** command.

```
edgeswitch:admin_06> nsallshow
{
010e00 020000 03f001 04f002
4 Nx_Ports in the Fabric }

edgeswitch:admin_06> nscamshow
nscam show for remote switches:
Switch entry for 1
  state  rev    owner
  known   v520  0xfffc02
  Device list: count 1
    Type Pid    COS     PortName                NodeName
    N    010e00;      3;10:00:00:00:00:01:00:00;10:00:00:00:00:00:01:00;
        Fabric Port Name: 20:0e:00:60:69:e2:18:b6
        Permanent Port Name: 10:00:00:00:00:01:00:00
        Port Index: 14
        Share Area: No
        Device Shared in Other AD: No

Switch entry for 3
  state  rev    owner
  known   v410  0xfffc02
  Device list: count 1
    Type Pid    COS     PortName                NodeName
    N    03f001;    2,3;10:00:00:00:c9:44:54:04;20:00:00:00:c9:44:54:04;
        FC4s: FCP
        NodeSymb: [36] "Emulex LP9002 FV3.92A2  DV5-5.10A10 "
        Fabric Port Name: 50:00:51:e3:70:9a:3d:e8
        Permanent Port Name: 10:00:00:00:c9:44:54:04
        Port Index: na
        Share Area: No
        Device Shared in Other AD: No

Switch entry for 4
  state  rev    owner
  known   v410  0xfffc02
  Device list: count 1
    Type Pid    COS     PortName                NodeName
    N    04f002;      3;10:00:00:00:00:03:00:00;10:00:00:00:00:00:03:00;
        Fabric Port Name: 50:06:06:91:23:45:6a:13
        Permanent Port Name: 10:00:00:00:00:03:00:00
        Port Index: na
        Share Area: No
        Device Shared in Other AD: No
```

All of the devices from both LSANs should appear in the output. If the devices do not appear in the output, issue the **cfgShow** command to verify your zone configuration. Use the **cfgactvshow** command to display the zone configuration currently in effect. The following example illustrates the use of **cfgactvshow**.

```
switch:admin> cfgactvshow
Effective configuration:
 cfg:   test
 zone:  lsan_san
        10:00:00:00:00:03:00:00
        10:00:00:00:00:01:00:00
 zone:  lsan_test
        50:06:01:60:38:e0:0b:a4
        10:00:00:00:c9:44:54:04
```

7. Log into the FC router and run the **lsanZoneShow** command to verify FIDs and devices to be shared among LSANs.

Administering FICON Fabrics

Chapter 12

IBM Fibre Connections (FICON®) is an industry-standard, high-speed input/output (I/O) interface for mainframe connections to storage devices. Fabric OS supports *intermix mode* operations, in which FICON and Fibre Channel technology work together. For specific information about intermix mode and other aspects of FICON, refer to the IBM Redbook, *FICON® Implementation Guide* (SG24-6497-01).

This chapter contains the following topics:

- "Overview of Fabric OS Support for FICON" on page 12-1
- "Configuring Switches" on page 12-6
- "Preparing a Switch" on page 12-7
- "Setting a Unique Domain ID" on page 12-9
- "Displaying Information" on page 12-10
- "Swapping Ports" on page 12-11
- "Clearing the FICON Management Database" on page 12-11
- "Using FICON CUP" on page 12-12
- "Troubleshooting" on page 12-19
- "Backing Up FICON Files" on page 12-20
- "Recording Configuration Information" on page 12-22

Overview of Fabric OS Support for FICON

Fabric OS provides standard support for FICON single-switch operation.

Control Unit Port (CUP) operation requires a Brocade FICON CUP license.

The following Fabric OS *standard features* support FICON fabrics:

- Port swapping

  Redirects resources from a failed port to a healthy port without changing the FICON host configuration. Port swapping is available for both FICON and open system environments. Port swapping resolves situations in which the hardware has failed and the channel configurations cannot be changed quickly. Port swapping has minimal or no impact on other switch features.

- Insistent domain ID (IDID)

  Allows the switch to insist on a specific domain ID before joining a fabric. This feature guarantees that a switch operates only with its preassigned domain ID.

- The FICON MIB module

  Addresses link incident data for FICON hosts and devices connected to a switch. It supplements other MIBs used to manage switches and should be used in conjunction with those other MIBs. For more information, see the *Fabric OS MIB Reference*.

- Link incident detection, registration, and reporting

  Provide administrative and diagnostic information.

- Switch Connection Control (SCC) policy

  Includes switch binding security methods that prevent unauthorized devices from joining a fabric.

These *optional features* provide further support:

- The Secure Fabric OS optional license

  Includes fabric binding, switch binding, and port binding security methods that prevent unauthorized devices from joining a fabric.

- A Fabric Manager optional license

  Can be used to manage a fabric that supports FICON and FCP devices and traffic. This is the recommended GUI management tool for FICON environments.

- Web Tools

  Can be used to manage a director (switch) that supports FICON and Fibre Channel Protocol (FCP) devices and traffic.

To incorporate and manage FICON on a switch or fabric, your system must have Fabric OS v4.1.2 or later installed. If you are implementing FICON in a single-switch *noncascaded* environment, there are no additional software requirements. The Brocade Advanced Zoning optionally licensed feature is required on all switches participating in a FICON multiple-switch *cascaded* environment.

Note

Some licenses are installed and activated on the switch at the factory. Use a Brocade management interface to verify that the required licenses are installed and activated on the switch.

The optional Secure Fabric OS license provides the following fabric, switch, and port binding features:

- *Fabric binding* is a security method for restricting switches within a multiple-switch fabric. The (SCC) policy prevents unauthorized switches from joining a fabric. Switches are authenticated using digital certificates and unique private keys provided to the Switch Link Authentication Protocol (SLAP).
- *Switch binding* is a security method for restricting devices that connect to a particular switch. If the device is another switch, this is handled by the SCC policy. If the device is a host or storage device, the Device Connection Control (DCC) policy binds those devices to a particular switch. Policies range from completely restrictive to reasonably flexible, based upon customer needs.
- *Port binding* is a security method for restricting host or storage devices that connect to particular switch ports. The DCC policy also binds device ports to switch ports. Policies range from completely restrictive to reasonably flexible, based upon customer needs.

For switches running Fabric OS 5.2.0 and higher, the SCC ACL with strict fabric-wide consistency can also be used for switch binding, in addition to the Secure Fabric OS mechanism.

Supported Switches

FICON protocol is supported on the following SilkWorm models and Fabric OS releases:

- SilkWorm 3900, Fabric OS v4.1.2 or later.
- SilkWorm 4100, Fabric OS v5.0.1b or later.
- SilkWorm 24000, Fabric OS v4.2.0 or later. The default one-domain configuration is supported; dual domain configurations and mixed FC4-16 port blade configurations are not supported.
- SilkWorm 48000, Fabric OS v5.0.1b or later. A single-domain configuration is supported with a mix of 16-port and 32-port blades. Dual-domain configurations are not supported. Mixed port blade configurations of SilkWorm 24000 and SilkWorm 48000 port blades (FC2-16, FC4-16, or FC4-32) in the same director are not supported in a FICON environment.

  The following port blades can exist in a FICON environment; however, FICON device connection to ports on these blades is not supported:

  - FC4-16IP
  - FC4-48
  - FR4-18i

  In an Admin Domain-enabled fabric, you should put all of the ports on these blades in an Admin Domain other than the one used for FICON ports. The ports on these blades should not belong to the zone in which FICON devices are present.

  The FC4-48 port blade is not supported for connecting to System z environments via FICON channels or via FCP zLinux on System z. To connect to the System z environment with the SilkWorm 48000, use the FC4-16 or FC4-32 port blades.

Both the SilkWorm 48000 and 4100 require the port-based routing policy either in a single switch configuration, or a cascaded switch configuration on switches in the fabric that have FICON devices attached (option 1 of the **aptPolicy** command). Other switches in the fabric can use the default exchange-based routing policy (option 3 of the **aptPolicy** command) only when Open Systems devices are attached to those switches.

Types of FICON Configurations

There are two types of FICON configurations:

- A *single-switch* configuration (called *switched point-to-point*) requires that the channel be configured to use single-byte addressing. If the channel is set up for two-byte addressing, then the cascaded configuration setup applies. This type of configuration is described in “Configuring a Single Switch” on page 12-7.
- A *cascaded configuration* (known as a *high integrity fabric*) requires a list of authorized switches. This authorization feature (called *fabric binding*) is available through Brocade Secure Fabric OS. The fabric binding policy allows a predefined list of switches (domains) to exist in the fabric and prevents other switches from joining the fabric. This type of configuration is described in “Configuring a High-Integrity Fabric” on page 12-7.

Control Unit Port (CUP)

Control Unit Port (CUP) protocol is used by IBM mainframe management programs to provide in-band management for FICON switches. When it is enabled, you can set up directors in a FICON environment to be managed through IBM mainframe management programs. CUP is an optionally licensed feature available with Fabric OS v4.4.0 or later.

CUP is supported on the SilkWorm 3900 and 24000 models running Fabric OS v4.4.0 or later. CUP is also supported on the SilkWorm 4100 and 48000 models running Fabric OS v5.0.1b or later.

For Fabric OS 5.2.0 and higher, the following restrictions apply to the SilkWorm 48000 when FICON Management Server mode (fmsmode) is enabled and CUP protocol is used to manage the switch:

- The switch is advertised to the mainframe via CUP as a 256-port switch (due to CUP protocol limitation).
- Port Information Block, PDCM, and port names are available for ports 0 through 254 only.
- CUP is not supported on the FC4-48 port blade Even though the FC4-48 port blade can be inserted in the chassis, only FCP ports can be attached to it.
- The FC4-48 port blade must not be inserted in slot 10 of the chassis. Slot 10 is reserved for the CUP port.
- For FR4-18i and FC4-16IP port blades, FICON channels and control units (CUs) can be attached only to the FC ports on these blades. These blades are advertised to the mainframe as 16-port and 8-port blades respectively.

FICON Commands

Table 12-1 summarizes the Fabric OS CLI commands that can be used for managing FICON fabrics. For detailed information on these commands, see the *Fabric OS Command Reference*.

Table 12-1 Fabric OS Commands Related to FICON and FICON CUP

Command	Description
Standard Fabric OS commands	
configure	Sets the domain ID and the insistent domain ID mode.
portSwap	Swaps ports.
portSwapDisable	Disables the portSwap command.
portSwapEnable	Enables the portSwap command.
portSwapShow	Displays information about swapped ports.
Commands specific to FICON	
ficonclear rlir	Removes all RLIR records from the local RLIR database.
ficonclear rnid	Removes all outdated RNID records from the local RNID database.
ficonshow ilir [fabric]	Displays FRU failure information on the local switch or on the fabric.
ficonshow lirr [fabric]	Displays registered listeners for link incidents for the local switch or for the fabric.

Table 12-1 Fabric OS Commands Related to FICON and FICON CUP (Continued)

Command	Description
ficonshow rlir [fabric]	Displays link incidents for the local switch or for the fabric.
ficonshow rnid [fabric]	Displays node identification data for all devices registered with the local switch or all devices registered with all switches defined in the fabric.
ficonshow switchrnid [fabric]	Displays node identification data for the local switch or for the fabric.
Commands specific to FICON CUP:	
ficoncupset fmsmode	Sets FICON Management Server mode on or off for the switch.
ficoncupset modereg	Sets the mode register bits for the switch.
ficoncupshow fmsmode	Displays the FICON Management Server mode setting for the switch.
ficoncupshow modereg	Displays the mode register bit settings for the switch.

Note

The Fabric OS CLI supports only a subset of the Brocade management features for FICON fabrics. The full set of FICON CUP administrative procedures is available using the Brocade Fabric Manager and Web Tools software features. You can also use an SNMP agent and the FICON Management Information Base (MIB). For information on these tools, see:

- Web Tools—*Web Tools Administrator's Guide*
- Fabric Manager—*Fabric Manager Administrator's Guide*
- SNMP Agent and FICON Management Information Base (MIB)—*Fabric OS MIB Reference*

Security Considerations

To administer FICON, you must have one of the following roles:

- Admin
- Operator
- SwitchAdmin
- FabricAdmin

The User and BasicSwitchAdmin roles are view-only. The ZoneAdmin and SecurityAdmin roles have no access.

In an Admin Domain-aware fabric, if you use the FICON commands (**ficonshow**, **ficonclear**, **ficoncupshow**, and **ficoncupset**) for any Admin Domain other than AD0 and AD255, the current switch must be a member of that Admin Domain. The output is not filtered based on the Admin Domain.

Configuring Switches

This section describes how to configure a switch in a FICON environment. Use the worksheet on page 12-22 to record your configuration information.

Following are recommended FICON environment configuration settings:

- Disable dynamic load sharing (**dlsReset** command).

  If DLS is enabled, traffic on existing ISL ports might be affected when one or more new ISLs is added between the same two switches. Specifically, adding the new ISL might result in dropped frames as routes are adjusted to take advantage of the bandwidth provided. By disabling DLS, you ensure that there will be no dropped frames.

  A similar situation occurs when an ISL port is taken offline and then brought back online. When the ISL port goes offline, the traffic on that port is rerouted to another ISL with a common destination. When the ISL port comes back online and DLS is enabled, the rerouting of traffic back to the ISL port might result in dropped frames. If DLS is not enabled, traffic will not be routed back.

- Configure ports that are connected to 1-Gbit/sec channels for fixed 1-Gbit/sec speed. Otherwise, when using fixed 1-Gbit/sec channels (both G5 and FICON Express), the FICON host might generate erroneous link incidents when the channels are coming online. These link incidents will result in a call home. Other than the generated link incident, the channel will come online and function normally.
- Enable in-order delivery (**iodSet** command).
- Enable VC translation link initialization on Extended Fabrics links, to stabilize them. See page 18-11 for details on this option of the **portCfgLongDistance** command.
- Although there are no specific zoning rules related to FICON environments, it is recommended that you follow standard FCP zoning practices. For management purposes, when operating in a mixed environment put FCP devices in one zone and FICON devices in another zone.
- The port-based routing policy is recommended for the SilkWorm 4100 and 48000 on any switch that has FICON devices attached. Other switches in the fabric with Open Systems devices exclusively can still use exchange-based routing.
- Some 1-Gbit/sec storage devices cannot auto-negotiate speed with the SilkWorm 48000 or 4100 ports. For these types of devices, configure ports that are connected to 1-Gbit/sec storage devices for fixed 1-Gbit/sec speed.

Preparing a Switch

To verify and prepare a switch for use in a FICON environment, complete the following steps:

1. Connect to the switch and log in as admin.
2. If not in a cascaded environment, proceed to step 3.

   If in a FICON cascaded environment, enter the following commands:

   - **licenseShow** to verify that required licenses (Secure Fabric OS and Zoning) are activated
   - **secModeShow** to determine if Secure Fabric OS is enabled; if it is disabled, enable it
   - **secPolicyShow** to verify that the SCC_POLICY is active
   - **pkiShow** to determine the existence of PKI objects, such as switch private key, private key passphrase, CSR, root certificate, and switch certificate. If none of these objects exists, see the *Secure Fabric OS Administrator's Guide* for information about creating the PKI objects and obtaining the digital certificate file.
3. Enter the **switchShow** command to verify that the switch and devices are online.
4. Change the routing policy on the switch from the default exchange-based policy to the required port-based policy for those switches with FICON devices directly attached. For the SilkWorm 4100, see the *Fabric OS Command Reference* for details about the **aptPolicy** command. For the SilkWorm 48000, see the *Web Tools Administrator's Guide*.
5. Enter the **ficonshow rnid** command to verify that the FICON devices are registered with the switch.
6. Enter the **ficonshow lirr** command to verify that the FICON host channels are registered to listen for link incidents.
7. Optionally, see "Using FICON CUP" on page 12-12for details about using FICON CUP.

Configuring a Single Switch

Single-switch configuration does not require IDID or fabric binding, provided that connected channels are configured for single-byte addressing. However, you should configure IDID to ensure that domain IDs are maintained.

Configuring a High-Integrity Fabric

To configure a high-integrity fabric (cascaded configuration):

1. Disable each switch in the fabric.
2. For each switch:
   a. Enable the IDID flag.
   b. Set the domain ID.
   c. Install security certificates and keys.
3. Enable the switches; this builds the fabric.

4. Set the SCC policy, as described in Chapter 4, “Configuring Standard Security Features”.
5. Set up security on the primary FCS switch.

   Use the **secModeEnable** command. The security policies are distributed to each switch in the fabric. (For details on the Quickmode procedure, see the *Secure Fabric OS Administrator’s Guide*.

Caution
If Security is enabled via the CLI in the FICON environment, then you should use the following syntax for the **secModeEnable** command:

```
secmodeenable --lockdown=scc --currentpwd --fcs "*"
```

Issuing the **secModeEnable** command as it appears above enables security and creates an SCC policy with all of the switches that currently reside in the fabric. It will also use the current password as the password for all available accounts on the switch.

Also, if you intend to use the **secModeEnable --quickmode** command, device connection control (DCC) policies are created for every port; this is not required for cascaded FICON configurations. These DCC policies in a cascaded configuration lock down the ports because no devices are logged into them before the command is executed, so no device can connect to these ports until you specifically add them to the DCC policy. If you issued the **secModeEnable --quickmode** command, the best solution is to delete the DCC policies that were created.

Be sure not to delete the SCC policy, which is required for FICON cascaded configurations.

6. Connect and enable channel and control unit (CU) devices. The Query for Security Attributes (QSA) response to the channel indicates that the fabric binding and IDID are enabled.

Figure 12-1 and Figure 12-2 show two viable cascaded configurations. These configurations require Channel A to be configured for two-byte addressing and require IDID and fabric binding. There can be only two switches in the path from the channel to the control unit.

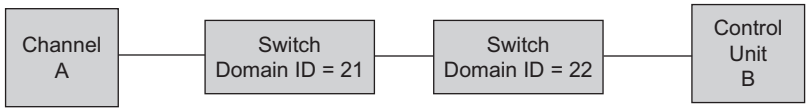

Figure 12-1 Cascaded Configuration, Two Switches

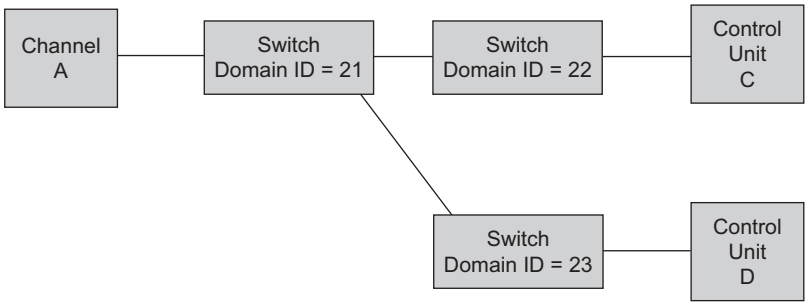

Figure 12-2 Cascaded Configuration, Three Switches

Setting a Unique Domain ID

In a cascaded configuration, each switch must have a unique domain ID, and insistent domain ID (IDID) mode must be enabled. To set a unique domain ID and enable IDID mode, complete the following steps:

1. Connect to the switch and log in as admin.
2. Verify that the switch has a unique domain ID. If it does not, set a unique domain ID.

   For instructions on displaying and changing the domain ID, see "Working With Domain IDs" on page 2-17.
3. Enter the **switchDisable** command to disable the switch.
4. Enter the **configure** command.
5. Enter **y** after the Fabric Parameters prompt.
6. To enable IDID mode, enter **y** after the "Insistent Domain ID Mode" prompt.

   (You can disable this mode by entering **n**.)
7. Respond to the remaining prompts (or press **Ctrl-d** to accept the other settings and exit).
8. Enter the **switchEnable** command to reenable the switch.

```
switch:admin> configure
Configure...
Fabric parameters (yes, y, no, n): [no] yes
Domain: (1..239) [3] 5
    R_A_TOV: (4000..120000) [10000]
    E_D_TOV: (1000..5000) [2000]
    Data field size: (256..2112) [2112]
    Sequence Level Switching: (0..1) [0]
    Disable Device Probing: (0..1) [0]
    Suppress Class F Traffic: (0..1) [0]
    VC Encoded Address Mode: (0..1) [0]
    Per-frame Route Priority: (0..1) [0]
    Long Distance Fabric: (0..1) [0]
    BB credit: (1..16) [16]
Insistent Domain ID Mode (yes, y, no, n): [no] y
Virtual Channel parameters (yes, y, no, n): [no]
Switch Operating Mode (yes, y, no, n): [no]
Zoning Operation parameters (yes, y, no, n): [no]
RSCN Transmission Mode (yes, y, no, n): [no]
Arbitrated Loop parameters (yes, y, no, n): [no]
System services (yes, y, no, n): [no]
Portlog events enable (yes, y, no, n): [no]
Committing configuration...done.
switch:admin>
```

Displaying Information

You can display link incidents, registered listeners, node identification data, and FRU failures, as described in the following procedures.

Link Incidents

The registered link incident record (RLIR) ELS contains the link incident information sent to a listener N_Port.

To display link incidents, connect to the switch, log in as user, and enter one of the following commands:

- For the local switch: **ficonshow rlir**
- For all switches defined in the fabric: **ficonshow rlir fabric**

Registered Listeners

To display registered listeners for link incidents, connect to the switch, log in as user, and enter one of the following commands:

- For the local switch: **ficonshow lirr**
- For all switches defined in the fabric: **ficonshow lirr fabric**

Node Identification Data

To display node-identification data, connect to the switch, log in as user, and enter any of the following commands:

- For the local switch: **ficonshow switchrnid**
- For all switches defined in the fabric: **ficonshow switchrnid fabric**
- For all devices registered with the local switch: **ficonshow rnid**
- For all devices registered with all switches defined in the fabric: **ficonshow rnid fabric**

FRU Failures

To display FRU failure information, connect to the switch, log in as admin, and enter one of the following commands:

- For the local switch: **ficonshow ilir**
- For all sw itches defined in the fabric: **ficonshow ilir fabric**

Swapping Ports

If a port malfunctions, or if you want to connect to different devices without having to re-wire your infrastructure, you can move a port's traffic to another port (*swap ports*) without changing the I/O Configuration Data Set (IOCDS) on the mainframe computer.

To swap ports, perform the following steps (see the example that follows):

1. Connect to the switch and log in as admin.
2. Enter the **portSwapEnable** command (to enable the command for port swapping).
3. Enter the **portDisable** command to disable the two ports to be swapped.
4. Enter the **portSwap** command to swap the ports.

   Any port in the switch can be used as the alternate for any other port within the same switch.
5. Reenable the ports using the **portEnable** command.
6. Enter **portSwapDisable** (to disable the command for port swapping).

In the following example:

```
switch:admin> portswapenable
switch:admin> portdisable [slot/] portA [slot/]portB
switch:admin> portswap [slot/] portA [slot/]portB
switch:admin> portenable [slot/] portA [slot/]portB
switch:admin> portswapdisable
```

- *slot* is the slot number of the port blade for a system with port blades (optional).
- *portA* is the original port number.
- *portB* is the alternate port number.

You can use the **portSwapShow** command to display information about swapped ports in a switch.

You can use the **portSwap** command to disable the portswap feature. You cannot use the **portSwap** command after this feature is disabled. The enabled state of the portswap feature is persistent across reboots and power cycles. Enabling and disabling the portswap feature does not affect previously executed portswap operations.

See the *Fabric OS Command Reference* for additional details about the **portSwap** command.

Clearing the FICON Management Database

You can clear RLIR and RNID records from the FICON management database as follows:

1. Connect to the switch and log in as admin.
2. To remove all the RLIR records from the local RLIR database, enter **ficonclear rlir**.
3. To remove all the RNID records marked "not current" from the local RNID database, enter **ficonclear rnid**.

Using FICON CUP

Host-based management programs manage switches using CUP protocol by sending commands to an emulated control device in Fabric OS. A Brocade switch that supports CUP can be controlled by one or more host-based management programs, as well as by Brocade tools.

A *mode register* controls the behavior of the switch with respect to CUP itself, and with respect to the behavior of other management interfaces.

FICON Management Server mode (fmsmode) must be enabled on the switch to enable CUP management features. When this mode is enabled, Fabric OS prevents local switch commands from interfering with host-based management commands by initiating *serialized access* to switch parameters.

If more than one switch is to be used in the FICON CUP fabric, Secure Fabric OS must be installed. See "Configuring a High-Integrity Fabric" on page 12-7 for more information.

If Brocade Advanced Zoning is in use, see "Zoning and PDCM Considerations" on page 12-18.

Setup Summary

To set up FICON CUP, use the following procedure and be sure to perform the steps in the order indicated.

To set up FICON CUP

1. Install Fabric OS v4.4.0 or later on a SilkWorm switch.
2. **For SilkWorm 24000 only**: Use the **portDisable** command to disable (block) port 126.

   For SilkWorm 48000 only: Use the **portDisable** command to disable (block) ports 254 and 255.

   Ports 126 (SilkWorm 12000) and 254 and 255 (SilkWorm 48000) are not supported in a CUP environment. After fmsmode has been successfully enabled, these two ports remain disabled and cannot be used either as an F_Port or an E_Port. Because these ports are not available after enabling fmsmode, you should first move any fiber connected to either of them to another free port.

   If you have an FR4-18i blade in your SilkWorm 48000, then the 16 virtual ports are disabled when you enable fmsmode.
3. Install a CUP license on the switch. See "Maintaining Licensed Features" on page 2-14.
4. Enable FICON management server mode (fmsmode) on the switch. See "Enabling and Disabling FICON Management Server Mode" on page 12-13.

After completing the setup, you can configure CUP attributes (FMS parameters). See "Setting Mode Register Bits" on page 12-16.

Note

After a nondisruptive firmware download from pre-5.2.0 firmware to 5.2.0 firmware (or higher), you must disable fmsmode and reenable it, so that FICON channels will query updated FICON parameters for CUP. See "Enabling and Disabling FICON Management Server Mode" on page 12-13.

Enabling and Disabling FICON Management Server Mode

To enable fmsmode:

1. Connect to the switch and log in as admin.
2. Enter **ficoncupset fmsmode enable**.

To disable fmsmode:

1. Connect to the switch and log in as admin.
2. Enter **ficoncupset fmsmode disable**.

The fmsmode setting can be changed whether the switch is offline or online. If fmsmode is changed while the switch is online, a device reset is performed for the control device and an RSCN is generated with PID 0xDDFE00 (where 0xDD is the domain ID of the switch).

When FICON Management Server mode is on, the Fabric OS CLI commands listed here return a "switch busy" response if they are issued when the host-based management tool is performing a write operation. This *serialization* prevents interference from local switch commands when a host-based management program is being used to administer the switch.

bladeDisable	bladeDisable
bladeEnable	bladeEnable
portDisable	switchCfgPersistentDisable
portEnable	switchDisable
portName	switchEnable
portShow	switchName
portSwap	switchShow

Note

You cannot use the **portCfgPersistentEnable** and **portCfgPersistentDisable** commands to persistently enable and disable ports when FICON Management Server mode is on. See the procedure "Persistently Enabling/Disabling Ports" on page 12-16.

Changing fmsmode from *disabled* to *enabled* triggers the following events:

- Access to switch parameters is serialized.
- The active CUP configuration data is established as follows:
  - Port and switch names are not read from the IPL; they remain as previously set.
  - Port "Block" and "Unblock" values are not read from the IPL; they remain as previously set with the **portEnable** and **portDisable** commands.
  - PDCM values are read from the IPL; the default is "Allow All."
- Brocade Advanced Zoning, if used, continues to be in force. If there are any differences in restrictions set up with Brocade Advanced Zoning and PDCM, the most restrictive rules are automatically applied.
- RSCNs are sent to devices if PDCM results in changes to connectivity between a set of ports.

Changing fmsmode from *enabled* to *disabled* triggers the following events:

1. A device reset is performed on the control device.
2. PDCM is no longer enforced.
3. RSCNs might be generated to some devices if PDCM removal results in changes to connectivity between a set of ports.
4. If a given port was set to “Block” or “Unblock,” that port remains disabled or enabled.
5. Serialized access to switch parameters ceases.

Setting Up CUP When FICON Management Server Mode Is Enabled

Fmsmode may be enabled and in use on a switch without a CUP License. The transition from fmsmode disabled to fmsmode enabled with the CUP license installed triggers the notification to the host systems that the CUP feature is available. Without this notification the host systems will never know the CUP feature is available, and consequently will never try to communicate with it. Hence, it is possible that fmsmode may already be enabled on the switch.

If FICON Management Server mode is already enabled, set up CUP as follows:

1. Verify that FICON Management Server mode is enabled by entering the **ficoncupshow fmsmode** command.

   If FICON Management Server mode is not enabled, see “Enabling and Disabling FICON Management Server Mode” on page 12-13.

   Caution
   If fmsmode is already enabled, disabling it might be disruptive to operation because ports that were previously prevented from communicating will now able to do so.

2. If FICON Management Server mode is enabled, then disable it by entering the **ficoncupset fmsmode disable** command.

   Install a CUP license key as described in “Adding and Removing FICON CUP Licenses” on page 12-17.

3. Enter the **ficoncupset fmsmode enable** command.

Displaying the fmsmode Setting

The **ficoncupshow fmsmode** command displays the effective fmsmode setting for the switch.

```
switch:admin> ficoncupshow fmsmode
fmsmode for the switch: Enabled
```

Displaying Mode Register Bit Settings

The mode register bits are described in Table 12-2.

Table 12-2 FICON CUP Mode Register Bits

POSC	Programmed offline state control. When this bit is set on, the host is prevented from taking the switch offline. The default setting is 1 (on).
UAM	User alert mode. When this bit is set on, a warning is issued when an action is attempted that will write CUP parameters on the switch. The default setting is 0 (off).
ASM	Active=saved mode. When this bit is set on, all CUP configuration parameters are persistent, meaning that they will be saved in nonvolatile storage in the initial program load (IPL) file that is applied upon a cold reboot or a power cycle. The default setting is 1 (on).
DCAM	Switch clock alert mode. When this bit is set on, a warning is issued when the **date**, **tsClockServer**, or **tsTimeZone** commands are entered to set the time and date on the switch. The default setting is 0 (off).
ACP	Alternate control prohibited. Because the Fabric OS CLI, Web Tools, and Fabric Manager are considered to be switch consoles, this bit has no effect on their operation. Attempts to set CUP parameters through SNMP are denied when this bit is set on. The default setting is 1 (on).
HCP	Host control prohibited. When this bit is set on, the host is not allowed to set CUP parameters. The default setting is 0 (off).

The **ficoncupshow modereg** command displays the mode register bit settings for the switch. A display of 0 indicates that the mode register bit is set to off; 1 indicates that the bit is set to on.

The command format is:

```
ficoncupshow modereg [bitname]
```

where *bitname* is one of the mode register bits described in Table 12-2.

To display all mode register bit settings for the switch:

```
switch:admin> ficoncupshow modereg
POSC  UAM  ASM  DCAM  ACP  HCP
------------------------------
   1    0    1     0    1    0
```

To display the mode register bit HCP for the switch:

```
switch:admin> ficoncupshow modereg HCP
HCP
0
```

Setting Mode Register Bits

Use the **ficoncupset modereg** command to set the FICON CUP mode register bits for the local switch. Consider the following when changing mode register bits:

- As required by the CUP protocol, the UAM bit cannot be changed using this command.
- All mode register bits except UAM are saved across power on/off cycles; the UAM bit is reset to 0 following a power-on.
- Mode register bits can be changed when the switch is offline or online. If the ACP or HCP bits are changed when the switch is online, they will take effect any time between the completion of the current command and the end of the CCW command chain (or the next alternate manager operation).

The command format is:

ficoncupset modereg [*bitname*] **0** | **1**

where:

bitname	One of the mode register bits described in Table 12-2 on page 12-15.
0	Specifies that the bit is off.
1	Specifies that the bit is on.

The following example sets the mode register bit HCP to off:

```
switch:admin> ficoncupset modereg HCP 0
Mode register bit HCP has been set to 0.
```

The following example sets the mode register bit ACP to on:

```
switch:admin> ficoncupset modereg ACP 1
Mode register bit ACP has been set to 1.
```

Persistently Enabling/Disabling Ports

When fmsmode is enabled, you cannot use the **portCfgPersistentEnable** and **portCfgPersistentDisable** commands to persistently enable and disable ports. Instead, use the following procedure:

1. Enter the **ficoncupshow modereg** command to display the mode register bit settings.
2. Verify that the ASM bit is set on (1).
3. If the ASM bit is set off (0), enter the **ficoncupset modereg asm 1** command to set it on

4. Use the **portEnable** and **portDisable** commands to enable and disable ports as necessary.

   The ports remain enabled or disabled after a switch reboot.

In the following example, the ASM bit is set to on; then, the port at slot 1, port 1 is enabled persistently:

```
switch:admin> ficoncupshow modereg
POSC  UAM  ASM  DCAM  ACP  HCP
------------------------------
   1    0    0     0    1    1

switch:admin> ficoncupset modereg ASM 1
Mode register bit ASM has been set to 1.

switch:admin> portenable 1/1
```

Port and Switch Naming Standards

Fabric OS handles differences in port and switch naming rules between CUP and itself as follows:

- CUP employs 8-bit characters in port address names and switch names; Fabric OS employs 7-bit characters. When fmsmode is enabled, all characters greater than 0x40 and not equal to 0xFF (EBCIDC code page 37 [0x25]) are allowed in the name; therefore, it is possible for a channel to set a name with nonprintable characters. If a name contains nonprintable characters, they are displayed as dots (...). The following characters are also displayed as dots: semicolon (;), comma (,), “equal” sign (=), and “at” sign (@).

  Note

  Configuration files that contain nonprintable characters should not be edited manually, because many editors replace nonprintable characters with some other characters without warning the user first.

- CUP has a 24-character unique port name limitation; Fabric OS supports port names up to 32 characters long. When fmsmode is enabled, names longer than 24 characters are truncated.
- To ensure that they are unique, the characters ~00, ~01, ~02, and so on are appended to port names.
- CUP allows a 24-character switch name; Fabric OS limits the switch name to 15 characters. To reconcile this difference, Fabric OS files the first 15 characters in the WWN record and stores the extra characters for CUP use.

Adding and Removing FICON CUP Licenses

If fmsmode is enabled when the FICON CUP license is removed, the control device is reset. PDCM enforcement continues. If fmsmode is disabled when the FICON CUP license is removed, no special action is taken.

If fmsmode is enabled on a switch that does not have a FICON CUP license and then the license is installed, you must first disable and then reenable fmsmode. If fmsmode is disabled and a FICON CUP license is installed, no special action is required.

Zoning and PDCM Considerations

The FICON Prohibit Dynamic Connectivity Mask (PDCM) controls whether or not communication between a pair of ports in the switch is prohibited or allowed. If there are any differences in restrictions set up with Brocade Advanced Zoning and PDCM, the most restrictive rules are automatically applied.

All FICON devices should be configured in a single zone using the “Domain, Area” notation. PDCM can then be used to “Allow” or “Prohibit” access between specific port pairs.

PDCM persists across a failover because it is replicated at all times to the standby CP blade. The active PDCM configuration is saved to the IPL if the ASM bit is set on.

Zoning and Link Incident Reporting

Non-implicit link incidents (such as NOS recognized or bit error rate threshold exceeded) that are related to a specific port are reported to registered listeners only if they are in the same zone as the port for which the link incident is being reported. As long as all FICON devices are in a single zone, link incidents for FICON ports are reported only to registered FICON listener ports.

The only exception to this is the *loss of synchronization* link incident. Even though this link incident is related to a specific port, it can be reported to a registered listener that is in a different zone.

Implicit link incidents (such as WWN card failure or power supply failure) that are not related to a specific port are reported to registered listeners regardless of zoning configuration.

Backing Up and Restoring Configurations

The Fabric OS **configUpload** command saves up to 16 FICON configuration files, including IPL files. For details on the behavior of the **configDownload** command, see “Restoring Configurations in a FICON Environment” on page 6-9.

Troubleshooting

The following sources provide useful problem-solving information:

- The standard support commands (**portLogDump**, **supportSave, supportShow**) or the Fabric Manager Event Log

  By default, the FICON group in the supportShow output is disabled. To enable the capture of FICON data in the supportShow output, enter the **supportshowcfgenable ficon** command. After you get confirmation that the configuration has been updated, the following will be collected and appear in the output for the supportShow command:

  - ficoncupshow fmsmode
  - ficoncupshow modereg
  - ficondbg dump rnid
  - ficondbg log
  - ficonshow ilir
  - ficonshow lirr
  - ficonshow rlir
  - ficonshow rnid
  - ficonshow switchrnid
  - ficucmd dump -A
- Other detailed information for protocol-specific problems:
  - Display port data structures using the **ptDataShow** command.
  - Display port registers using the **ptRegShow** command.

Identifying Ports

The **ficonshow rlir** command displays, among other information, a tag field for the switch port. You can use this tag to identify the port on which a FICON link incident occurred. The tag field is a concatenation of the switch domain ID and port number, in hexadecimal format. The following example shows a link incident for the switch port at domain ID 120, port 93 (785d in hex):

```
switch:admin> ficonshow rlir
{
 {Fmt  Type PID    Port   Incident Count  TS Format   Time Stamp
  0x18 F    785d00   93               1  Time server Thu Apr 22 09:13:32 2004
  Port Status:          Link not operational
  Link Failure Type:    Loss of signal or synchronization

  Registered Port WWN     Registered Node WWN      Flag  Node Parameters
  50:05:07:64:01:40:16:03 50:05:07:64:00:c1:69:ca  0x10  0x200115
  Type number:          002064
  Model number:         103
  Manufacturer:         IBM
  Plant of Manufacture: 02
  Sequence Number:      0000000169CA
  tag:                  155d

  Switch Port WWN         Switch Node WWN          Flag  Node Parameters
  20:5d:00:60:69:80:45:7c 10:00:00:60:69:80:45:7c  0x00  0x200a5d
  Type number:          SLKWRM
  Model number:         24K
  Manufacturer:         BRD
  Plant of Manufacture: CA
  Sequence Number:      000000000078
  tag:                  785d
}
}
The Local RLIR database has 1 entry.
switch:admin> ficonshow rlir
```

Backing Up FICON Files

The FICON file access facility is used to store configuration files. This includes IPL and other configuration files. The Fabric OS saves the IPL and all other configuration files on the switch. A maximum of 16 configuration files, including the IPL file, are supported.

You can upload the configuration files saved on the switch to a management workstation using the **configUpload** command. If the switch loses the configuration due to a hardware failure or filesystem error, use the **configDownload** command to restore previously uploaded configuration files. Because data uploaded using the **configUpload** command also contains the IPL, if *Active=Saved* mode is enabled, then the switch will ignore the IPL file downloaded via the **configDownload** command.

To upload the configuration files

Enter the **configUpload** command.

When you execute the **configUpload** command, all the files saved in the file access facility are uploaded to a management workstation (there is a section in the uploaded configuration file labeled “FICON_CUP” that exists in an encoded format).

To download configuration files with Active=Saved mode enabled

Enter the **configDownload** command.

The contents of existing files saved on the switch, which are also present in the “FICON_CUP” section, are overwritten.

The files in the “FICON” section of the configuration file, which are not currently on the switch, are saved on the switch.

The IPL will not be replaced because *Active=Saved* mode is enabled. A warning message is displayed in the event log to warn users that the IPL will not be overwritten.

To download configuration files with Active=Saved mode disabled

Enter the **configDownload** command.

The contents of existing files saved on the switch, which are also present in the “FICON_CUP” section, are overwritten.

The files in the “FICON” section of configuration file, which are not currently on the switch, are saved on the switch.

The IPL will be replaced because *Active=Saved* mode is disabled.

Recording Configuration Information

You can use the following worksheet for recording FICON configuration information.

FICON® Switch Configuration Worksheet									
FICON® Switch Manufacturer:_____Type: _____ Model: _____ S/N: _____									
HCD Defined Switch ID_____(Switch ID) **FICON® Switch Domain ID_____(Switch @)**					**Cascaded Directors No _____Yes _____** **Corresponding Cascaded Switch Domain ID _____** **Fabric Name _____**				
FICON® Switch F_Ports					**Attached N_Ports / E_Ports (CU, CPC, or ISL)**				
Slot Number	Port Number	Port Address	Laser Type: LX / SX	Port Name	Node Type CU / CHNL	Machine Type	Model	Serial Number	ISL CU I/F CPC CHPID

Sample IOCP Configuration File for SilkWorm 3900, 4100, 24000, and 48000 Switches

The channel subsystem controls communication between a configured channel, the control unit, and the device. The I/O Configuration Dataset (IOCDS) defines the channels, control units, and devices to the designated logical partitions (LPARs) within the server; this is defined using the Input/Output Configuration Program (IOCP). The IOCP statements are typically built using the hardware configuration dialog (HCD). The interactive dialog is used to generate your Input/Output Definition File (IODF), invoke the IOCP program, and subsequently build your production IOCDS.

Each FICON director in a fabric must have a unique domain ID and a unique switch ID. The switch ID used in the IOCP definitions can be any value between x'00' to x'FF'. The domain ID range for directors is hex x'01' to x'EF' or decimal 1 to 239. When defining the switch IDs in the IOCP definitions, ensure that you use values within the FICON director's range.

The switch ID has to be assigned by the user and must be unique within the scope of the definitions (IOCP and HCD).

The domain ID is assigned by the manufacturer and can be customized to a different value. It must be unique within the fabric.

Brocade recommends that the switch ID (in IOCP or HCD) be set to the same value as the domain ID of the FICON director, which is defined to the FICON director at installation time. This simplifies the configuration and reduces confusion by having a common designation across all definitions.

For more information on switch numbering, refer to the IBM publication *FICON® Implementation Guide* (SG24-6497-00).

In the following sample IOCP configuration file, the UNIT value for FICON CUP definitions is **2032** for any FICON director regardless of vendor or platform. So all SilkWorm 3900, SilkWorm 4100, or SilkWorm 24000 switches require **UNIT=2032** for the CUP definition. All Domain IDs are specified in hex values in the IOCP (and not in decimal values); the Domain IDs in the example are for demonstration purposes only.

```
*--------------------------------------------------------------------
* SilkWorm 24000 Domain_ID=61 (in hex)
*--------------------------------------------------------------------
 CNTLUNIT CUNUMBR=0D8,UNITADD=00,UNIT=2032,
          PATH=(50,51),
               LINK=(61FE,61FE)
 IODEVICE ADDRESS=(0D8,1),CUNUMBR=0D8,UNIT=2032,STADET=Y,UNITADD=00
*--------------------------------------------------------------
 CNTLUNIT CUNUMBR=0D9,UNITADD=00,UNIT=2032,
               PATH=(8A,8B),
               LINK=(22FE,22FE)
 IODEVICE ADDRESS=(0D9,1),CUNUMBR=0D9,UNIT=2032,STADET=Y,UNITADD=00
*--------------------------------------------------------------------
* SilkWorm 3900 Domain_ID=25 (in hex)
*--------------------------------------------------------------------
*
 CNTLUNIT CUNUMBR=0DB,UNITADD=00,UNIT=2032,
               PATH=(5A,5B),
               LINK=(25FE,25FE)
 IODEVICE ADDRESS=(0DB,1),CUNUMBR=0DB,UNIT=2032,STADET=Y,UNITADD=00
*
*--------------------------------------------------------------------
```

Sample RMF Configuration File for Mainframe

Resource management facility (RMF) is a performance management tool that measures selected areas of system activity, including different views of the FICON channel. RMF presents data collected in the form of System Management Facility (SMF) records. This data is essential for any kind of FICON channel performance troubleshooting.

To obtain an RMF FICON director activity report, you must include the keyword “FCD” in the RMF configuration file for the FICON director (this is generic for any FICON director). You must also define the CUP port. In the sample below, the keyword is boldfaced.

```
/*********************************************************************/
/*  MONITOR I OPTIONS                                                */
/*                                                                   */
/*                 X A     O N L Y                                   */
/*                                                                   */
/*********************************************************************/
         FCD                       /* FICON Director    */
         CHAN                      /* COLLECT CHANNEL STATISTICS    */
         CPU                       /* COLLECT CPU STATISTICS        */
         CYCLE(1000)               /* SAMPLE ONCE EVERY SECOND      */
         DEVICE(NOSG)              /* PREVENT SORT OF STORAGE GROUPS*/
         DEVICE(NOCHRDR)           /* CHARACTER READER STATISTICS
                                      WILL NOT BE COLLECTED         */
         DEVICE(COMM)              /* COMMUNICATION EQUIPMENT STATS.
                                      WILL BE COLLECTED PDS 1/25/94 */
         DEVICE(DASD)              /* DIRECT ACCESS DEVICE STATISTICS
                                      WILL BE COLLECTED             */
         DEVICE(GRAPH)             /* GRAPHICS DEVICE STATISTICS
                                      WILL BE COLLECTED             */
         DEVICE(TAPE)              /* TAPE DEVICE STATISTICS
                                      WILL BE COLLECTED             */
         DEVICE(NOUNITR)           /* UNIT RECORD DEVICE STATISTICS
                                      WILL NOT BE COLLECTED         */
         DEVICE(NONMBR)            /* NO DEVICE SELECTIVITY BY
                                      DEVICE NUMBERS                */
         IOQ(DASD)                 /* COLLECT DASD I/O QUEUING
                                      STATISTICS                    */
         IOQ(NOCHRDR)              /* PREVENT CHARACTER READER
                                      I/O QUEUING STATISTICS        */
         IOQ(NOUNITR)              /* PREVENT UNIT RECORD DEVICE
                                      I/O QUEUING STATISTICS        */
         IOQ(COMM)                 /* COLLECT COMMUNICATION EQUIPMENT
                                      I/O QUEUING STATS PDS 1/25/94 */
         IOQ(GRAPH)                /* COLLECT GRAPHICS DEVICE
```

Configuring the Distributed Management Server

Chapter 13

The Brocade Fabric OS Distributed Management Server allows a SAN management application to retrieve information and administer interconnected switches, servers, and storage devices. The management server assists in the autodiscovery of switch-based fabrics and their associated topologies.

This chapter contains the following topics:

- "Enabling and Disabling the Platform Services" on page 13-1
- "Controlling Access" on page 13-2
- "Configuring the Server Database" on page 13-6
- "Controlling Topology Discovery" on page 13-7

A client of the management server can find basic information about the switches in the fabric and use this information to construct topology relationships. The management server also allows you to obtain certain switch attributes and, in some cases, modify them. For example, logical names identifying switches can be registered with the management server.

The management server provides several advantages for managing a Fibre Channel fabric:

- It is accessed by an external Fibre Channel node at the well-known address *FFFFFAh*, so an application can access information about the entire fabric management with minimal knowledge of the existing configuration.
- It is replicated on every SilkWorm switch within a fabric.
- It provides an unzoned view of the overall fabric configuration. This fabric topology view exposes the internal configuration of a fabric for management purposes; it contains interconnect information about switches and devices connected to the fabric. Under normal circumstances, a device (typically an FCP initiator) queries the Name Server for storage devices within its member zones. Because this limited view is not always sufficient, the management server provides the application with a list of the entire Name Server database.

Note

Management server platform service is available only with Fabric OS v2.3.0 and later.

Enabling and Disabling the Platform Services

The management server is located at the Fibre Channel well-known address *FFFFFAh*. All management services except platform services are enabled by default.

Note

The commands **msplMgmtActivate** and **msplMgmtDeactivate** are allowed only in AD0 and AD255.

To enable platform services

1. Connect to the switch and log in as admin.
2. Enter the **msplMgmtActivate** command.

```
switch:admin> msplmgmtactivate
Request to activate MS Platform Service in progress......
*Completed activating MS Platform Service in the fabric!
switch:admin>
```

To disable platform services

1. Connect to the switch and log in as admin.
2. Enter the **msplMgmtActivate** command.
3. Enter **y** to confirm the deactivation.

```
switch:admin> msplmgmtdeactivate
MS Platform Service is currently enabled.
This will erase MS Platform Service configuration
information as well as database in the entire fabric.
Would you like to continue this operation? (yes, y, no, n): [no] y
Request to deactivate MS Platform Service in progress......
*Completed deactivating MS Platform Service in the fabric!
switch:admin>
```

Controlling Access

You can use the **msConfigure** command to control access to the management server database.

An access control list (ACL) of WWN addresses determines which systems have access to the management server database. The ACL typically contains those WWNs of host systems that are running management applications.

If the list is empty (the default), the management server is accessible to all systems connected in-band to the fabric. For more access security, you can specify WWNs in the ACL so that access to the management server is restricted to only those WWNs listed.

The ACL is switch-based. Therefore, only hosts that are connected directly to the switch are affected by the ACL. A host that is somewhere else in the fabric and is connected to a switch with an empty ACL is allowed to access the management server.

Note

The **msConfigure** command is disabled if the switch is in secure mode. See the *Secure Fabric OS Administrator's Guide* for more information.

To display the management server ACL

1. Connect to the switch and log in as admin.
2. Enter the **msConfigure** command.

   The command becomes interactive.

3. At the select prompt, enter **1** to display the access list.

   A list of WWNs that have access to the management server is displayed.

   In the following example, the list is empty:

```
switch:admin> msconfigure
0       Done
1       Display the access list
2       Add member based on its Port/Node WWN
3       Delete member based on its Port/Node WWN
select : (0..3) [1] 1
MS Access list is empty.
0       Done
1       Display the access list
2       Add member based on its Port/Node WWN
3       Delete member based on its Port/Node WWN
done ...
switch:admin>
```

To add a member to the ACL

1. Connect to the switch and log in as admin.
2. Enter the **msConfigure** command.

   The command becomes interactive.
3. At the select prompt, enter **2** to add a member based on its port/node WWN.
4. Enter the WWN of the host to be added to the ACL.
5. At the prompt, enter **1** to verify the WWN you entered was added to the ACL.
6. After verifying that the WWN was added correctly, enter **0** at the prompt to end the session.
7. At the “Update the FLASH?” prompt, enter **y**.

8. Press **Enter** to update the nonvolatile memory and end the session.

```
switch:admin> msconfigure
0       Done
1       Display the access list
2       Add member based on its Port/Node WWN
3       Delete member based on its Port/Node WWN
select : (0..3) [1] 2
Port/Node WWN (in hex): [00:00:00:00:00:00:00:00] 20:00:00:20:37:65:ce:aa
*WWN is successfully added to the MS ACL.
0       Done
1       Display the access list
2       Add member based on its Port/Node WWN
3       Delete member based on its Port/Node WWN
select : (0..3) [2] 1
MS Access List consists of (14): {
 20:00:00:20:37:65:ce:aa
 20:00:00:20:37:65:ce:bb
 20:00:00:20:37:65:ce:ff
 20:00:00:20:37:65:ce:11
 20:00:00:20:37:65:ce:22
 20:00:00:20:37:65:ce:33
 20:00:00:20:37:65:ce:44
 10:00:00:60:69:04:11:24
 10:00:00:60:69:04:11:23
 21:00:00:e0:8b:04:70:3b
 10:00:00:60:69:04:11:33
 20:00:00:20:37:65:ce:55
 20:00:00:20:37:65:ce:66
 00:00:00:00:00:00:00:00
}
0       Done
1       Display the access list
2       Add member based on its Port/Node WWN
3       Delete member based on its Port/Node WWN
select : (0..3) [1] 0
done ...
Update the FLASH?  (yes, y, no, n): [yes] y
*Successfully saved the MS ACL to the flash.
switch:admin>
```

To delete a member from the ACL

1. Connect to the switch and log in as admin.
2. Enter the **msConfigure** command.

   The command becomes interactive.
3. At the select prompt, enter **3** to delete a member based on its port/node WWN.
4. At the prompt, enter the WWN of the member to be deleted from the ACL.
5. At the prompt, enter **1** to verify the WWN you entered was deleted from the ACL.
6. After verifying that the WWN was deleted correctly, enter **0** at the prompt to end the session.
7. At the “Update the FLASH?” prompt, enter **y**.

8. Press **Enter** to update the nonvolatile memory and end the session.

```
switch:admin> msconfigure
0       Done
1       Display the access list
2       Add member based on its Port/Node WWN
3       Delete member based on its Port/Node WWN
select : (0..3) [1] 3
Port/Node WWN (in hex): [00:00:00:00:00:00:00:00] 20:00:00:20:37:65:ce:aa
*WWN is successfully deleted from the MS ACL.
0       Done
1       Display the access list
2       Add member based on its Port/Node WWN
3       Delete member based on its Port/Node WWN
select : (0..3) [2] 1
MS Access List consists of (13): {
 20:00:00:20:37:65:ce:aa
 20:00:00:20:37:65:ce:bb
 20:00:00:20:37:65:ce:ff
 20:00:00:20:37:65:ce:11
 20:00:00:20:37:65:ce:22
 20:00:00:20:37:65:ce:33
 10:00:00:60:69:04:11:24
 10:00:00:60:69:04:11:23
 21:00:00:e0:8b:04:70:3b
 10:00:00:60:69:04:11:33
 20:00:00:20:37:65:ce:55
 20:00:00:20:37:65:ce:66
}
0       Done
1       Display the access list
2       Add member based on its Port/Node WWN
3       Delete member based on its Port/Node WWN
select : (0..3) [1] 0
done ...
Update the FLASH?  (yes, y, no, n): [yes] y
*Successfully saved the MS ACL to the flash.
switch:admin>
```

Configuring the Server Database

The management server database can be viewed or cleared. The command **msPlClearDB** will be allowed only in AD0 and AD255.

To view the contents of the management server database

1. Connect to the switch and log in as admin.
2. Enter the **msPlatShow** command.

   The contents of the management server platform database are displayed.

```
switch:admin> msplatshow
---------------------------------------------------------
Platform Name: [9] "first obj"
Platform Type: 5 : GATEWAY
Number of Associated M.A.: 1
[35] "http://java.sun.com/products/plugin"
Number of Associated Node Names: 1
Associated Node Names:
10:00:00:60:69:20:15:71
---------------------------------------------------------
Platform Name: [10] "second obj"
Platform Type: 7 : HOST_BUS_ADAPTER
Number of Associated M.A.: 1
Associated Management Addresses:
[30] "http://java.sun.com/products/1"
Number of Associated Node Names: 1
Associated Node Names:
10:00:00:60:69:20:15:75
```

To clear the management server database

1. Connect to the switch and log in as admin.
2. Enter the **msplClearDb** command.
3. Enter **y** to confirm the deletion.

   The management server platform database is cleared.

Controlling Topology Discovery

The topology discovery feature can be displayed, enabled, and disabled; it is disabled by default. The commands **mstdEnable and mstdDisable** are allowed only in AD0 and AD255.

To display topology discovery status

1. Connect to the switch and log in as admin.
2. Enter the **mstdReadConfig** command.

```
switch:admin> mstdreadconfig
*MS Topology Discovery is Enabled.
switch:admin>
```

To enable topology discovery

1. Connect to the switch and log in as admin.
2. Enter the **mstdEnable** command to enable the discovery feature locally.
3. Enter the **mstdEnable all** command to enable the discovery feature on the entire fabric.

```
switch:admin> mstdenable

Request to enable MS Topology Discovery Service in progress....
*MS Topology Discovery enabled locally.

switch:admin> mstdenable ALL

Request to enable MS Topology Discovery Service in progress....
*MS Topology Discovery enabled locally.
*MS Topology Discovery Enable Operation Complete!!
```

To disable topology discovery

1. Connect to the switch and log in as admin.
2. Enter the **mstdDisable** command to disable the discovery feature locally.

   A warning displays that all NID entries might be cleared.
3. Enter **y** to disable the discovery feature.
4. Enter the **mstdDisable all** command to disable the discovery feature on the entire fabric.

5. Enter **y** to disable the discovery feature.

Note

Disabling management server topology discover might erase all NID entries.

```
switch:admin> mstddisable
This may erase all NID entries. Are you sure?  (yes, y, no, n): [no] y

Request to disable MS Topology Discovery Service in progress....
*MS Topology Discovery disabled locally.

switch:admin> mstddisable all
This may erase all NID entries. Are you sure?  (yes, y, no, n): [no] y

Request to disable MS Topology Discovery Service in progress....
*MS Topology Discovery disabled locally.
*MS Topology Discovery Disable Operation Complete!!
```

Working With Diagnostic Features

Chapter 14

This chapter provides information on diagnostics and how to display system, port, and specific hardware information. It also describes how to set up system logging mapping (**syslogd**) and how to set up the offloading of error messages (**supportSave**).

This chapter contains the following topics:

- "About Fabric OS Diagnostics" on page 14-1
- "Viewing Switch Status" on page 14-4
- "Viewing Port Information" on page 14-6
- "Viewing Equipment Status" on page 14-9
- "Viewing the System Message Log" on page 14-11
- "Viewing the Port Log" on page 14-11
- "Configuring for syslogd" on page 14-13
- "Viewing and Saving Diagnostic Information" on page 14-15
- "Setting Up Automatic Trace Dump Transfers" on page 14-16

About Fabric OS Diagnostics

The purpose of the diagnostic subsystem is to evaluate the integrity of the system hardware.

Diagnostics are invoked two ways:

- Automatically during the power-on self test (POST)
- Manually using Fabric OS CLI commands

The error messages generated during these test activities are sent to the serial console and system message logs, whose output formats may differ slightly.

Use the **diagHelp** command to receive a list of all available diagnostic commands.

See the *Fabric OS Command Reference* for a complete description of each command.

Viewing Power-On Self Test

By default, when you power on the system, the boot loader automatically performs power-on self tests and loads a Fabric OS kernel image.

The POST tests provide a quick indication of hardware readiness when hardware is powered up. These tests do not require user input to function. They typically operate within several minutes, and support minimal validation because of the restriction on test duration. Their purpose is to give a basic health check before a new switch joins a fabric.

These tests are divided into two groups: POST1 and POST2. POST1 validates the hardware interconnect of the device, and POST2 validates the ability of the device to pass data frames between the ports. The specific set of diagnostic and test commands run during POST depends on the switch model.

The factory default configuration is set to run POST2, but you can configure your switch to bypass POST2, which runs after the kernel image has started but before general system services such as login are enabled.

Although each test performed during POST2 is configurable, you should only modify a POST2 test if directed by your switch provider's customer service representative.

You can use the **diagDisablePost** command to disable both POST1 and POST2, and you can reenable it using the **diagEnablePost** command. See the *Fabric OS Command Reference* for additional information about these commands.

The following example shows a typical boot sequence, including POST messages:

```
The system is coming up, please wait...

Read board ID of 0x80 from addr 0x23
Read extended model ID of 0x16 from addr 0x22
Matched board/model ID to platform index 4
PCI Bus scan at bus 0
:   :   :
:   :   :
Checking system RAM - press any key to stop test

Checking memory address: 00100000

System RAM test using Default POST RAM Test succeeded.

Press escape within 4 seconds to enter boot interface.
Booting "Fabric Operating System" image.

Linux/PPC load:
BootROM command line: quiet
Uncompressing Linux...done.
Now booting the kernel
Attempting to find a root file system on hda2...
modprobe: modprobe: Can't open dependencies file /lib/modules/2.4.19/modules.dep (No such file or directory)
INIT: version 2.78 booting
INIT: Entering runlevel: 3
eth0: Link status change: Link Up. 100 Mbps Full duplex Auto (autonegotiation complete).

INITCP: CPLD Vers: 0x95 Image ID: 0x19
uptime: 2008; sysc_qid: 0
Fabric OS (Paulsa45)
Paulsa45 console login: 2005/03/31-20:12:42, [TRCE-5000], 0,, INFO, ?, trace:, trace_buffer.c, line: 1170

2005/03/31-20:12:42, [LOG-5000], 0,, INFO, SW4100_P45, Previous message repeat 1 time(s), trace_ulib.c, line: 540
2005/03/31-20:12:43, [HAM-1004], 219,, INFO, SW4100_P45, Processor rebooted - Unknown
SNMP Research SNMP Agent Resident Module Version 15.3.1.4
Copyright 1989, 1990, 1991, 1992, 1993, 1994, 1995, 1996, 1997, 1998, 1999, 2000, 2001 SNMP Research, Inc.
sysctrld: all services Standby
FSSK 2: chassis0(0): state not synchronized
FSSK 2: Services starting a COLD recovery
2005/03/31-20:12:48, [FSS-5002], 0,, INFO, SW4100_P45, chassis0(0): state not synchronized, svc.c, line: 318
2005/03/31-20:12:48, [FSS-5002], 0,, INFO, SW4100_P45, Services starting a COLD recovery, mdev.c, line: 638
2005/03/31-20:12:49, [MFIC-1002], 220,, INFO, Paulsa45, Chassis FRU header not programmed for switch NID, using
defaults (applies only to FICON environments).
sysctrld: all services Active
2005/03/31-20:12:50, [DGD-5001], 0,, INFO, SW4100_P45, Slot 0 has started POST., main.c, line: 1189
POST1: Started running Thu Mar 31 20:12:51 GMT 2005
POST1: Test #1 - Running turboramtest
POST1: Test #2 - Running portregtest
POST1: Script PASSED with exit status of 0 Thu Mar 31 20:12:54 GMT 2005 took (0:0:3)
POST2: Started running Thu Mar 31 20:12:55 GMT 2005
POST2: Test #1 - Running portloopbacktest (SERDES)
POST2: Test #2 - Running minicycle (SERDES)
POST2: Running diagshow
POST2: Script PASSED with exit status of 0 Thu Mar 31 20:13:12 GMT 2005 took (0:0:17)
2005/03/31-20:13:13, [BL-1000], 221,, INFO, Paulsa45, Initializing Ports... Enabling switch...
2005/03/31-20:13:13, [BL-1001], 222,, INFO, Paulsa45, Port Initialization Completed
2005/03/31-20:13:13, [EM-5012], 0,, INFO, SW4100_P45, EM: sent dumpready to ME., em.c, line: 2152
2005/03/31-20:13:13, [DGD-5002], 0,, INFO, SW4100_P45, Slot 0 has passed the POST tests., main.c, line: 936
```

If you choose to bypass POST2, or after POST2 completes, various system services are started and the boot process displays additional console status and progress messages.

Viewing Switch Status

Use the **switchStatusShow** command to display the overall status of the switch, including its power supplies, fans, and temperature. If the status of any one of these components is either marginal or down, the overall status of the switch is also displayed as marginal or down. If all components have a healthy status, the switch displays a healthy status.

To modify the rules used to classify the health of each component use the **switchStatusPolicySet** command. To view the rules, use the **switchStatusPolicyShow** command.

To view the overall status of the switch

1. Connect to the switch and log in as admin.
2. Enter the **switchStatusShow** command:

```
switch:admin> switchstatusshow
Switch Health Report                        Report time: 03/21/2005 03:50:36 PM
Switch Name:    SWFCR
IP address:     10.33.54.176
SwitchState:    MARGINAL
Duration:       863:23
Power supplies monitor  MARGINAL
Temperatures monitor    HEALTHY
Fans monitor            HEALTHY
WWN servers monitor     HEALTHY
Standby CP monitor      HEALTHY
Blades monitor          HEALTHY
Flash monitor           HEALTHY
Marginal ports monitor  HEALTHY
Faulty ports monitor    HEALTHY
Missing SFPs monitor    HEALTHY
All ports are healthy
switch:admin>
```

For more information on how the overall switch status is determined, see the **switchStatusPolicySet** command in the *Fabric OS Command Reference*.

To display switch information

1. Connect to the switch and log in as admin.
2. At the command line, enter the **switchShow** command. This command displays the following information for a switch:
   - switchname - Displays the switch name.
   - switchtype - Displays the switch model and firmware version numbers.
   - switchstate - Displays the switch state: Online, Offline, Testing, or Faulty.
   - switchrole - Displays the switch role: Principal, Subordinate, or Disabled.
   - switchdomain - Displays the switch Domain ID.
   - switchid - Displays the embedded port D_ID of the switch.
   - switchwwn - Displays the switch World Wide Name.
   - switchbeacon - Displays the switch beaconing state: either ON or OFF.

The **switchShow** command also displays the following information for ports on the specified switch:

- Module type - The SFP type if a SFP is present.
- Port speed - The speed of the Port (1G, 2G, 4G, N1, N2, N4, or AN). The speed can be fixed, negotiated, or auto negotiated.
- Port state - The port status.
- Comment - Displays information about the port. This section might be blank or display WWN for F_Port or E_Port, Trunking state, upstream or downstream status.

The details displayed for each switch differ on different switch models. For more information see the **switchShow** command in the *Fabric OS Command Reference*.

To display the uptime for a switch

1. Connect to the switch and log in as admin.
2. At the command line, enter the **uptime** command:

```
switch:admin> uptime
4:43am  up 1 day, 12:32,  1 user,  load average: 1.29, 1.31, 1.27
switch:admin>
```

The **uptime** command displays the length of time the system has been in operation, the total cumulative amount of uptime since the system was first powered-on, the date and time of the last reboot (applies only to Fabric OS v3.x and v2.6.x systems), the reason for the last reboot (applies only to Fabric OS v3.x and v2.6.x systems), and the load average over the past one minute (1.29 in the preceding example), five minutes (1.31 in the example), and 15 minutes (1.27 in the example). The reason for the last switch reboot is also recorded in the system message log.

Viewing Port Information

Use the commands that follow to view information about ports.

To view the status of a port

1. Connect to the switch and log in as admin.
2. Enter the **portShow** command, specifying the number that corresponds to the port you are troubleshooting. In this example, the status of port two is shown:

```
switch:admin> portshow 2
portName:
portHealth: HEALTHY

Authentication: None
portDisableReason: None
portCFlags: 0x1
portFlags: 0x4903        PRESENT ACTIVE E_PORT G_PORT U_PORT LOGICAL_ONLINE LOGIN LED
portType:  10.0
portState: 1    Online
portPhys:  6    In_Sync
portScn:   16   E_Port    Trunk port
port generation number:    351
portId:    290200
portIfId:    43020004
portWwn:   20:02:00:05:1e:34:01:be
portWwn of device(s) connected:
        None
Distance:  normal
portSpeed: N4Gbps

Interrupts:        0          Link_failure: 50         Frjt:         0
Unknown:           0          Loss_of_sync: 55         Fbsy:         0
Lli:               524        Loss_of_sig:  54
Proc_rqrd:         0          Protocol_err: 0
Timed_out:         0          Invalid_word: 0
Rx_flushed:        0          Invalid_crc:  0
Tx_unavail:        0          Delim_err:    0
Free_buffer:       0          Address_err:  0
Overrun:           0          Lr_in:        100
Suspended:         0          Lr_out:       50
Parity_err:        0          Ols_in:       50
2_parity_err:      0          Ols_out:      52
CMI_bus_err:       0

switch:admin>
```

See the *Fabric OS Command Reference* for additional **portShow** command information, such as the syntax for slot or port numbering.

To display the port statistics

1. Connect to the switch and log in as admin.
2. At the command line, enter the **portStatsShow** command.

Port statistics include information such as number of frames received, number of frames sent, number of encoding errors received, and number of class 2 and class 3 frames received.

See the *Fabric OS Command Reference* for additional **portStatsShow** command information, such as the syntax for slot or port numbering.

```
switch:admin> portstatsshow 3/7
stat_wtx                 0           4-byte words transmitted
stat_wrx                 0           4-byte words received
stat_ftx                 0           Frames transmitted
stat_frx                 0           Frames received
stat_c2_frx              0           Class 2 frames received
stat_c3_frx              0           Class 3 frames received
stat_lc_rx               0           Link control frames received
stat_mc_rx               0           Multicast frames received
stat_mc_to               0           Multicast timeouts
stat_mc_tx               0           Multicast frames transmitted
tim_rdy_pri              0           Time R_RDY high priority
tim_txcrd_z              0           Time BB credit zero
er_enc_in                0           Encoding errors inside of frames
er_crc                   0           Frames with CRC errors
er_trunc                 0           Frames shorter than minimum
er_toolong               0           Frames longer than maximum
er_bad_eof               0           Frames with bad end-of-frame
er_enc_out               0           Encoding error outside of frames
er_bad_os                0           Invalid ordered set
er_c3_timeout            0           Class 3 frames discarded due to timeout
er_c3_dest_unreach       0           Class 3 frames discarded due to destination unreachable
er_other_discard         0           Other discards
er_crc_good_eof          0           Crc error with good eof
er_inv_arb               0           Invalid ARB
open                     0           loop_open
transfer                 0           loop_transfer
opened                   0           FL_Port opened
starve_stop              0           tenancies stopped due to starvation
fl_tenancy               0           number of times FL has the tenancy
nl_tenancy               0           number of times NL has the tenancy
zero_tenancy             0           zero tenancy

switch:admin>
```

To display a summary of port errors for a switch

1. Connect to the switch and log in as admin.
2. At the command line, enter the **portErrShow** command. See the *Fabric OS Command Reference* for additional portErrShow command information.

```
switch:admin> porterrshow
       frames  enc  crc  too  too  bad  enc disc link loss loss frjt fbsy
     tx   rx   in  err shrt long  eof  out   c3 fail sync sig
sig=================================================================
 0:   22   24    0    0    0    0    0  1.5m    0    7    3    0    0    0
 1:   22   24    0    0    0    0    0  1.2m    0    7    3    0    0    0
 2:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
 3:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
 4:  149m  99m    0    0    0    0    0  448    0    7    6    0    0    0
 5:  149m  99m    0    0    0    0    0  395    0    7    6    0    0    0
 6:  147m  99m    0    0    0    0    0  706    0    7    6    0    0    0
 7:  150m  99m    0    0    0    0    0  160    0    7    5    0    0    0
 8:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
 9:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
10:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
11:    0    0    0    0    0    0    0    0    0    0    0    2    0    0
12:    0    0    0    0    0    0    0    0    0    0    0    2    0    0
13:    0    0    0    0    0    0    0    0    0    0    0    2    0    0
14:    0    0    0    0    0    0    0    0    0    0    0    2    0    0
15:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
32:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
33:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
34:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
35:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
36:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
37:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
38:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
39:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
40:   99m 146m    0    0    0    0    0  666    0    6  796    7    0    0
41:   99m 149m    0    0    0    0    0   15k   0    2  303    4    0    0
42:   99m 152m    0    0    0    0    0  665    0    2  221    5    0    0
43:   99m 147m    0    0    0    0    0   16k   0    2  144    4    0    0
44:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
45:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
46:    0    0    0    0    0    0    0    0    0    0    0    2    0    0
47:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
```

The **portErrShow** command output provides one output line per port. See Table 14-1 for a description of the error types.

Table 14-1 Error Summary Description

Error Type	Description
frames tx	Frames transmitted
frames rx	Frames received
enc in	Encoding errors inside frames
crc err	Frames with CRC errors
too shrt	Frames shorter than minimum
too long	Frames longer than maximum
bad eof	Frames with bad end-of-frame delimiters
enc out	Encoding error outside of frames
disc c3	Class 3 frames discarded
link fail	Link failures (LF1 or LF2 states)

Table 14-1 Error Summary Description (Continued)

Error Type	Description
loss sync	Loss of synchronization
loss sig	Loss of signal
frjt	Frames rejected with F_RJT
fbsy	Frames busied with F_BSY

Viewing Equipment Status

You can display status for fans, power supply, and temperature.

Note

The number of fans, power supply units, and temperature sensors depends on the switch type. For detailed specifications on these components, refer to the switch hardware reference manual.

The specific output from the status commands varies depending on the switch type.

To display the status of the fans

1. Connect to the switch and log in as admin.
2. Enter the **fanShow** command:

```
switch:admin> fanshow
Fan 1 is OK          speed is 7010 RPM
Fan 2 is OK          speed is 7180 RPM
Fan 3 is OK          speed is 7068 RPM
Fan 4 is OK          speed is 7116 RPM
Fan 5 is OK          speed is 7155 RPM
Fan 6 is OK          speed is 7001 RPM
switch:admin>
```

The possible status values are:

OK—Fan is functioning correctly.

Absent—Fan is not present.

Below minimum—Fan is present but rotating too slowly or stopped.

Above minimum—Fan is rotating too quickly.

Unknown—Unknown fan unit installed.

FAULTY—Fan has exceeded hardware tolerance.

To display the status of a power supply

1. Connect to the switch and log in as admin.
2. Enter the **psShow** command:

```
switch:admin> psshow
Power Supply #1 is OK
0335,FF2Z0007161,60-0000739-02, B,,DCJ3002-01P, B,FF2Z0007161
Power Supply #2 is faulty
0335,FF2Z0007176,60-0000739-02, B,,DCJ3002-01P, B,FF2Z0007176
switch:admin>
```

The possible status values are:

OK—Power supply functioning correctly.

Absent—Power supply not present.

Unknown—Unknown power supply unit installed.

Predicting failure—Power supply is present but predicting failure.

FAULTY—Power supply is present but faulty (no power cable, power switch turned off, fuse blown, or other internal error).

To display temperature status

1. Connect to the switch and log in as admin.
2. At the command line, enter the **tempShow** command:

```
switch:admin> tempshow

Index       Status       Centigrade       Fahrenheit
---------------------------------------------------
 1           OK             21               70
 2           OK             22               72
 3           OK             29               84
 4           OK             24               75
 5           OK             25               77
switch:admin>
```

Information displays for each temperature sensor in the switch.

The possible temperature status values are:

OK—Temperature is within acceptable range.

FAIL—Temperature is outside of acceptable range. Damage might occur.

Viewing the System Message Log

The system message log feature enables messages to be saved across power cycles and reboots.

SilkWorm 12000, 24000, and 48000 director models maintain an independent system message log for each of the two CP blades. For these models, you should configure syslogd to support chronological system message logs. For details, see "Configuring for syslogd" on page 14-13.

For details on error messages, see the *Fabric OS Message Reference*.

To display the system message log, with no page breaks

1. Connect to the switch and log in as admin.
2. Enter the **errDump** command at the command line.

To display the system message log one at a time

1. Connect to the switch and log in as admin.
2. Enter the **errShow** command at the command line.

To clear the system message log

1. Connect to the switch and log in as admin.
2. Enter the **errClear** command at the command line.

   All switch and chassis events are removed.

Viewing the Port Log

The Fabric OS maintains an internal log of all port activity. The port log stores entries for each port as a circular buffer. Each port has space to store 8000 log entries. When the log is full, the newest log entries overwrite the oldest log entries. Port logs are not persistent and are lost over power-cycles and reboots. If the port log is disabled, an error message displays.

Note

Port log functionality is completely separate from the system message log. Port logs are typically used to troubleshoot device connections.

To view the port log

1. Connect to the switch and log in as admin.
2. Enter the **portLogShow** command:

```
switch:admin> portlogshow 12

time          task       event  port cmd  args
-----------------------------------------------
Thu Apr 14 12:07:09 2005
12:07:09.350  PORT       Rx        0   40  02fffffd,00fffffd,0608ffff,14000000
12:07:09.350  PORT       Tx        0    0  c0fffffd,00fffffd,060807fc
12:07:10.812  PORT       Tx        0   40  02fffffd,00fffffd,07feffff,14000000
12:07:10.813  PORT       Rx        0    0  c0fffffd,00fffffd,07fe0627
12:07:19.492  PORT       Tx        4   40  02fffffd,00fffffd,0800ffff,14000000
12:07:19.492  PORT       Tx       22   40  02fffffd,00fffffd,0802ffff,14000000
12:07:19.493  PORT       Rx        4    0  c0fffffd,00fffffd,08009287
12:07:19.493  PORT       Tx       24   40  02fffffd,00fffffd,0804ffff,14000000
12:07:19.494  PORT       Tx       31   40  02fffffd,00fffffd,0806ffff,14000000
12:07:19.494  PORT       Rx       22    0  c0fffffd,00fffffd,0802928d
12:07:19.494  PORT       Rx       24    0  c0fffffd,00fffffd,080492a3
12:07:19.495  PORT       Rx       31    0  c0fffffd,00fffffd,080692a7
```

Use the commands summarized in Table 14-2 to view and manage port logs.

Table 14-2 Commands for Port Log Management

Command	Description
portLogClear	Clear port logs for all or particular ports.
portLogDisable	Disable port logs for all or particular ports.
portLogDump	Display port logs for all or particular ports, without page breaks.
portLogEnable	Enable port logs for all or particular ports.
portLogShow	Display port logs for all or particular ports, with page breaks.
Note: See the *Fabric OS Command Reference* for additional information about these commands.	

The **portLogDump** command output (trace) is a powerful tool that is used to troubleshoot fabric issues. The **portLogDump** output provides detailed information about the actions and communications within a fabric. By understanding the processes that are taking place in the fabric, issues can be identified and located.

The **portLogDump** command displays the port log, showing a portion of the Fibre Channel payload and header (FC-PH). The header contains control and addressing information associated with the frame. The payload contains the information being transported by the frame and is determined by the higher-level service or FC_4 upper level protocol. There are many different payload formats based on the protocol.

Because a **portLogDump** output is long, a truncated example is presented:

```
switch:admin> portlogdump
task event port cmd args
-----------------------------------------------
16:30:41.780 PORT Rx 9 40 02ffffd,00ffffd,0061ffff,14000000
16:30:41.780 PORT Tx 9 0 c0ffffd,00ffffd,0061030f
16:30:42.503 PORT Tx 9 40 02ffffd,00ffffd,0310ffff,14000000
16:30:42.505 PORT Rx 9 0 c0ffffd,00ffffd,03100062
16:31:00.464 PORT Rx 9 20 02fffc01,00fffca0,0063ffff,01000000
16:31:00.464 PORT Tx 9 0 c0fffca0,00fffc01,00630311
16:31:00.465 nsd ctin 9 fc 000104a0,0000007f
16:31:00.465 nsd ctout 9 fc 00038002,00000003,01fffc01
16:31:00.466 PORT Tx 9 356 03fffca0,00fffc01,00630311,01000000
16:31:00.474 PORT Rx 9 0 c0fffc01,00fffca0,00630311
16:31:01.844 PORT Tx 9 40 02ffffd,00ffffd,0312ffff,14000000
16:31:01.854 PORT Rx 9 0 c0ffffd,00ffffd,03120064
16:31:01.963 PORT Rx 9 40 02ffffd,00ffffd,0065ffff,14000000
16:31:01.963 PORT Tx 9 0 c0ffffd,00ffffd,00650313
16:31:14.726 INTR pstate 0 LF2
16:31:14.729 PORT scn 0 137 00000000,00000000,00000008
16:31:14.729 PORT scn 0 129 00000000,00000000,00000400
16:31:14.729 PORT scn 0 2 00010004,00000000,00000002
16:31:14.730 SPEE sn 0 ws 00000002,00000000,00000000
<output truncated>
```

Configuring for syslogd

The system logging daemon (syslogd) is a process on UNIX, Linux, and some Windows systems that reads and logs messages as specified by the system administrator.

Fabric OS can be configured to use a UNIX-style syslogd process to forward system events and error messages to log files on a remote host system.

The host system can be running UNIX, Linux, or any other operating system that supports the standard syslogd functionality.

Note

Fabric OS releases earlier than v4.4.0 did not support UNIX local7 facilities; they supported kern facilities.

Starting with Fabric OS v4.4.0, kern facilities are no longer supported; UNIX local7 facilities are supported (the default facility level is 7).

Configuring for syslogd involves configuring the host, enabling syslogd on the SilkWorm model, and, optionally, setting the facility level.

Configuring the Host

Fabric OS supports a subset of UNIX-style message severities that default to the UNIX local7 facility. To configure the host, edit the */etc/syslog.conf* file to map Fabric OS message severities to UNIX severities, as shown in Table 14-3. In

Table 14-3 Fabric OS to UNIX Message Severities

Fabric OS Message Severity	UNIX Message Severity
Critical (1)	Emergency (0)
Error (2)	Error (3)
Warning (3)	Warning (4)
Info (4)	Info (6)

In this example, Fabric OS messages map to local7 facility level 7 in the */etc/syslog.conf* file:

```
local7.emerg          /var/adm/swcritical
local7.alert          /var/adm/alert7
local7.crit           /var/adm/crit7
local7.err            /var/adm/swerror
local7.warning        /var/adm/swwarning
local7.notice         /var/adm/notice7
local7.info           /var/adm/swinfo
local7.debug          /var/adm/debug7
```

If you prefer to map Fabric OS severities to a different UNIX local7 facility level, see "To set the facility level" on page 14-15.

Configuring the Switch

Configuring the switch involves specifying syslogd hosts and, optionally, setting the facility level. You can also remove a host from the list of syslogd hosts.

To specify syslogd hosts

1. Connect to the switch and log in as admin.
2. Enter the **syslogDipAdd** command and specify an IP address.
3. Verify that the IP address was entered correctly, using the **syslogDipShow** command.

You can specify up to six host IP addresses for storing syslog messages, as shown in this example:

```
switch:admin> syslogdipadd 10.1.2.1
switch:admin> syslogdipadd 10.1.2.2
switch:admin> syslogdipadd 10.1.2.3
switch:admin> syslogdipadd 10.1.2.4
switch:admin> syslogdipadd 10.1.2.5
switch:admin> syslogdipadd 10.1.2.6
switch:admin> syslogdipshow
syslog.IP.address.1 10.1.2.1
syslog.IP.address.2 10.1.2.2
syslog.IP.address.3 10.1.2.3
syslog.IP.address.4 10.1.2.4
syslog.IP.address.5 10.1.2.5
syslog.IP.address.6 10.1.2.6
```

To set the facility level

1. Connect to the switch and log in as admin.
2. Enter the following command:

```
switch:admin> syslogdfacility -l n
```

   n is a number from 0 through 7, indicating a UNIX local7 facility. The default is 7.

It is necessary to set the facility level only if you specified a facility other than local7 in the host */etc/syslog.conf* file.

To remove a syslogd host from the list

1. Connect to the switch and log in as admin.
2. Enter the **syslogDipRemove** command:

```
switch:admin> syslogdipremove 10.1.2.1
```

3. Verify the IP address was deleted using the **syslogDipShow** command.

Viewing and Saving Diagnostic Information

Enter the **supportShow** command to dump important diagnostic and status information to the session screen, where you can review it or capture its data.

To save a set of files that customer support technicians can use to further diagnose the switch condition, enter the **supportSave** command. The command prompts for an FTP server, packages the following files, and sends them to the specified server:

- The output of the **supportShow** command
- The contents of any trace dump files on the switch
- System message logs (for SilkWorm directors, **supportSave** saves the system message logs from both of the CP blades)

See also to "Setting Up Automatic Trace Dump Transfers," next.

Setting Up Automatic Trace Dump Transfers

You can set up a switch so that diagnostic information is transferred automatically to a remote server. Then, if a problem occurs you can provide your customer support representative with the most detailed information possible. To ensure the best service, you should set up for automatic transfer as part of standard switch configuration, before a problem occurs.

Setting up for automatic transfer of diagnostic files involves the following tasks:

- Specify a remote server to store the files.
- Enable the automatic transfer of trace dumps to the server. (Trace dumps overwrite each other by default; sending them to a server preserves information that would otherwise be lost.)
- You should also set up a periodic checking of the remote server so that you are alerted if the server becomes unavailable and you can correct the problem.

After the setup is complete, you can run the **supportSave -c** command to save diagnostic information to the server (without specifying server details).

The following procedures describe in detail the tasks for setting up automatic transfer. For details on the commands, see the *Fabric OS Command Reference*.

To specify a remote server

1. Verify that the FTP service is running on the remote server.
2. Connect to the switch and log in as admin.
3. Enter the following command:

```
switch:admin> supportftp -s
```

   The command is interactive; provide the required information.

4. Respond to the prompts as follows:

Host Name	Enter the name or IP address of the server where the file is to be stored; for example, 192.1.2.3.
User name	Enter the user name of your account on the server; for example, “JohnDoe”.
Password	Enter your account password for the server.
Remote directory	Specify a path name for the remote directory. Absolute path names can be specified using forward slash (/). Relative path names create the directory in the user’s home directory on UNIX servers, and in the directory where the FTP server is running on Windows servers.

To enable the automatic transfer of trace dumps

1. Connect to the switch and log in as admin.
2. Enter the following command:

```
switch:admin> traceftp -e
```

To set up periodic checking of the remote server

1. Connect to the switch and log in as admin.
2. Enter the following command:

```
switch:admin> supportftp -t interval
```

The interval is in hours. The minimum interval is 1 hour. Specify 0 hours to disable the checking feature.

To save a comprehensive set of diagnostic files to the server

1. Connect to the switch and log in as admin.
2. Enter the following command:

```
switch:admin> supportsave -c
```

Troubleshooting

Chapter 15

This chapter provides information on troubleshooting and the most common procedures used to diagnose and recover from problems. It also includes specific troubleshooting scenarios as examples.

This chapter contains the following topics:

- "About Troubleshooting" on page 15-1
- "Most Common Problem Areas" on page 15-2
- "Gathering Information for Technical Support" on page 15-3
- "Analyzing Connection Problems" on page 15-5
- "Restoring a Segmented Fabric" on page 15-8
- "Correcting Zoning Setup Issues" on page 15-10
- "Recognizing MQ-WRITE Errors" on page 15-12
- "Correcting I2C Bus Errors" on page 15-13
- "Correcting Device Login Issues" on page 15-14
- "Identifying Media-Related Issues" on page 15-19
- "Correcting Link Failures" on page 15-21
- "Correcting Marginal Links" on page 15-24
- "Inaccurate Information in the System Message Log" on page 15-25
- "Recognizing the Port Initialization and FCP Auto Discovery Process" on page 15-26
- "Port Mirroring" on page 15-27

About Troubleshooting

Troubleshooting should begin at the center of the SAN — the fabric. Because switches are located between the hosts and storage devices and have visibility into both sides of the storage network, starting with them can help narrow the search path. After eliminating the possibility of a fault within the fabric, see if the problem is on the storage side or the host side, and continue a more detailed diagnosis from there. Using this approach can quickly pinpoint and isolate problems.

For example, if a host cannot detect a storage device, run a switch command (such as **switchShow**) to find out if the storage device is logically connected to the switch. If not, focus first on the switch directly connecting to storage. Use your vendor-supplied storage diagnostic tools to better understand why it is not visible to the switch. If the storage can be detected by the switch, and the host still cannot detect the storage device, then there is still a problem between the host and switch.

Most Common Problem Areas

See Table 15-1 for a list of the most common problem areas that arise within SANs and a list of tools that can be used to resolve them.

Table 15-1 Common Troubleshooting Problems and Tools

Problem Area	Investigate	Tools
Fabric	• Missing devices • Marginal links (unstable connections) • Incorrect zoning configurations • Incorrect switch configurations	• Switch LEDs • Switch commands (for example, switchShow or nsAllShow) for diagnostics • Web or GUI-based monitoring and management software tools
Storage Devices	• Physical issues between switch and devices • Incorrect storage software configurations	• Device LEDs • Storage diagnostic tools • Switch commands (for example, switchShow or nsAllShow) for diagnostics
Hosts	• Downlevel HBA firmware • Incorrect device driver installation • Incorrect device driver configuration	• Use the latest HBA firmware recommended by the switch supplier or on the HBA supplier's web site • Host operating system diagnostic tools • Device driver diagnostic tools • Switch commands (for example, switchShow or nsAllShow) for diagnostics
Storage Management Applications	Incorrect installation and configuration of the storage devices that the software references. For example, if using a volume-management application, check for: • Incorrect volume installation • Incorrect volume configuration	Application-specific tools and resources

Gathering Information for Technical Support

If you are troubleshooting a production system, you need to gather data quickly. As soon as a problem is observed, perform the following tasks (if using a dual CP system, run the commands on both CPs):

1. Enter the **supportSave** command to save RASLOG, TRACE, and supportShow (active CP only) information for the local CP to a remote FTP location. On a dual CP system, only the local CP information is saved and supportShow information is not available on the active CP.

   Note

   It is recommended that you use the **supportFtp** command to set up the supportSave environment for automatic dump transfers and the use of the **-n -c** options; this will save you from having to enter (or know) all the required FTP parameters needed to successfully execute a supportSave operation.

   For details about the supportSave command, see "Viewing and Saving Diagnostic Information" on page 14-15 and "Setting Up Automatic Trace Dump Transfers" on page 14-16

2. Enter the **pdShow** command to display data from a panic dump file. The panic dump file contains information that might be useful to determine the cause of the system panic.
3. Enter the **saveCore** command to save or remove core files created by daemons.

For more details about these commands, see the *Fabric OS Command Reference*.

Troubleshooting Questions

Common steps and questions to ask yourself when troubleshooting a system problem are as follows:

1. What is the current Fabric OS level?
2. What is the switch hardware version?
3. Is the switch operational?
4. Impact assessment and urgency:
   - Is the switch down?
   - Is it a standalone switch?
   - Are there VE, VEX or EX ports connected to the SilkWorm 48000?
   - How large is the fabric?
   - Is it a secure fabric?
   - Is the fabric redundant?
5. Run the **supportSave** command. (See "Viewing and Saving Diagnostic Information" on page 14-15 and "Setting Up Automatic Trace Dump Transfers" on page 14-16).
6. Document the sequence of events by answering the following questions:
   - What happened just prior to the problem?
   - Is the problem reproducible?
   - If so, what are the steps to produce the problem?
   - What configuration was in place when the problem occurred?

7. Did a failover occur?
8. Was security enabled?
9. Was POST enabled?
10. Are serial port (console) logs available?
11. Which CP blade was active? (only applicable to the SilkWorm 24000 and SilkWorm 48000)
12. What and when were the last actions or changes made to the system environment?

Use the following steps to retrieve as much of the following informational items as possible prior to contacting the SAN technical support vendor.

1. Switch information:
    - Serial number (located on the chassis)
    - World Wide Name (obtain using **licenseIdShow** or **wwn** commands)
    - Fabric OS version (obtain using the **version** command)
    - Switch configuration settings
    - **supportSave** output
    - **pdShow** and **saveCore** output
2. Host information:
    - OS version and patch level
    - HBA type
    - HBA firmware version
    - HBA driver version
    - Configuration settings
3. Storage information:
    - Disk/tape type
    - Disk/tape firmware level
    - Controller type
    - Controller firmware level
    - Configuration settings
    - Storage software (such as EMC Control Center, Veritas SPC, etc.)

Analyzing Connection Problems

If a host is unable to detect its target (for example, a storage or tape device), you should begin troubleshooting the problem in the middle of the data path. Determine if the problem is *above* or *below* the starting point, then continue to divide the suspected problem path in half until you can pinpoint the problem.

To check the logical connection

1. Enter the **switchShow** command.
2. Review the output and determine if the device is logically connected to the switch:
   - A device that *is* logically connected to the switch will be registered as an F_Port or L_Port.
   - A device that is *not* logically connected to the switch will be registered as something other than an F_Port or L_Port.
3. If the missing device *is* logically connected, proceed to the next troubleshooting procedure ("To check the Name Server (NS)" on page 15-7.
4. If the missing device is *not* logically connected, check the device and everything on that side of the data path. Also see "Correcting Link Failures" on page 15-21.

   This includes all aspects of the host OS, the driver settings and binaries, the device Basic Input Output System (BIOS) settings, the SFP, the cable going from the switch to the device, the SFP on the switch side of that cable, and all switch settings related to the device. See "To check for a loop initialization failure" on page 15-22 as the next potential trouble spot.

To check for Fibre Channel connectivity problems

1. Enter the **fcPing** command, which:
   - Checks the zoning configuration for the two ports specified.
   - Generates an ELS (Extended Link Service frame) ECHO request to the source port specified and validates the response.
   - Generates an ELS ECHO request to the destination port specified and validates the response.
2. Regardless of the device's zoning, the **fcPing** command sends the ELS frame to the destination port. A device can take any one of the following actions:
   - Send an ELS Accept to the ELS request.
   - Send an ELS Reject to the ELS request.
   - Ignore the ELS request.

   There are some devices that do not support the ELS ECHO request. In these cases, the device will either not respond to the request or send an ELS reject. When a device does not respond to the ELS request, further debugging is required; however, do not assume that the device is not connected to the Fibre Channel.

Following is sample output from the **fcPing** command in which one device accepts the request and another device rejects the request:

```
switch:admin> fcping 10:00:00:00:c9:29:0e:c4 21:00:00:20:37:25:ad:05
Source:      10:00:00:00:c9:29:0e:c4
Destination: 21:00:00:20:37:25:ad:05
Zone Check:  Not Zoned

Pinging 10:00:00:00:c9:29:0e:c4 [0x20800] with 12 bytes of date:
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1162 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1013 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1442 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1052 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1012 usec
5 frames sent, 5 frames received, 0 frames rejected, 0 frames timeout
Round-trip min/avg/max = 1012/1136/1442 usec

Pinging 21:00:00:20:37:25:ad:05 [0x211e8] with 12 bytes of data:
Request rejected
Request rejected
Request rejected
Request rejected
Request rejected
5 frames sent, 0 frames received, 5 frames rejected, 0 frames timeout
Round-trip min/avg/max = 0/0/0 usec
switch:admin>
```

Following is sample output from the **fcPing** command in which one device accepts the request and another device does not respond to the request:

```
switch:admin> fcping 0x020800 22:00:00:04:cf:75:63:85
Source:      0x20800
Destination: 22:00:00:04:cf:75:63:85
Zone Check:  Zoned

Pinging 0x020800 with 12 bytes of data:
received reply from 0x020800: 12 bytes time:1159 usec
received reply from 0x020800: 12 bytes time:1006 usec
received reply from 0x020800: 12 bytes time:1008 usec
received reply from 0x020800: 12 bytes time:1038 usec
received reply from 0x020800: 12 bytes time:1010 usec
5 frames sent, 5 frames received, 0 frames rejected, 0 frames timeout
Round-trip min/avg/max = 1006/1044/1159 usec

Pinging 22:00:00:04:cf:75:63:85 [0x217d9] with 12 bytes of data:
Request timed out
Request timed out
Request timed out
Request timed out
Request timed out
5 frames sent, 0 frames received, 0 frames rejected, 5 frames timeout
Round-trip min/avg/max = 0/0/0 usec
switch:admin>
```

For details about the **fcPing** command, see the *Fabric OS Command Reference*.

To check the Name Server (NS)

1. Enter the **nsShow** command on the switch to which the device is attached:

```
The Local Name Server has 9 entries {

 Type Pid   COS     PortName              NodeName            TTL(sec)

*N   021a00;   2,3;20:00:00:e0:69:f0:07:c6;10:00:00:e0:69:f0:07:c6; 895
    Fabric Port Name: 20:0a:00:60:69:10:8d:fd
NL   051edc;    3;21:00:00:20:37:d9:77:96;20:00:00:20:37:d9:77:96; na
    FC4s: FCP [SEAGATE ST318304FC      0005]

    Fabric Port Name: 20:0e:00:60:69:10:9b:5b
 NL   051ee0;    3;21:00:00:20:37:d9:73:0f;20:00:00:20:37:d9:73:0f; na
    FC4s: FCP [SEAGATE ST318304FC      0005]

    Fabric Port Name: 20:0e:00:60:69:10:9b:5b
 NL   051ee1;    3;21:00:00:20:37:d9:76:b3;20:00:00:20:37:d9:76:b3; na
    FC4s: FCP [SEAGATE ST318304FC      0005]

    Fabric Port Name: 20:0e:00:60:69:10:9b:5b
 NL   051ee2;     3;21:00:00:20:37:d9:77:5a;20:00:00:20:37:d9:77:5a; na
    FC4s: FCP [SEAGATE ST318304FC      0005]

    Fabric Port Name: 20:0e:00:60:69:10:9b:5b
 NL   051ee4;    3;21:00:00:20:37:d9:74:d7;20:00:00:20:37:d9:74:d7; na
    FC4s: FCP [SEAGATE ST318304FC      0005]

    Fabric Port Name: 20:0e:00:60:69:10:9b:5b
 NL   051ee8;    3;21:00:00:20:37:d9:6f:eb;20:00:00:20:37:d9:6f:eb; na
    FC4s: FCP [SEAGATE ST318304FC      0005]

    Fabric Port Name: 20:0e:00:60:69:10:9b:5b
 NL   051eef;    3;21:00:00:20:37:d9:77:45;20:00:00:20:37:d9:77:45; na
    FC4s: FCP [SEAGATE ST318304FC      0005]

    Fabric Port Name: 20:0e:00:60:69:10:9b:5b
 N    051f00;   2,3;50:06:04:82:bc:01:9a:0c;50:06:04:82:bc:01:9a:0c; na
    FC4s: FCP [EMC     SYMMETRIX       5267]

    Fabric Port Name: 20:0f:00:60:69:10:9b:5b
```

2. Look for the device in the NS list, which lists the nodes connected to that switch, allowing you to determine if a particular node is accessible on the network.

   - If the device is *not* present in the NS, the problem is between the storage device and the switch. There might be a time-out communication problem between edge devices and the Name Server, or there might be a login issue. First check the edge device documentation to determine if there is a time-out setting or parameter that can be reconfigured. You should also check the port log for NS registration information and FCP probing failures (using the **fcpProbeShow** command). If these queries do not help solve the problem, contact the support organization for the product that appears to be inaccessible.
   - If the device *is* listed in the NS, the problem is between the storage device and the host. There might be a zoning mismatch or a host/storage issue. Proceed to "To check for zoning problems" on page 15-8.

3. Enter the **portLoginShow** command to check the port login status.
4. Enter the **fcpProbeShow** command to display the FCP probing information for the devices attached to the specified F_Port or L_Port. This information includes the number of successful logins and SCSI INQUIRY commands sent over this port and a list of the attached devices.
5. Check the port log to determine whether or not the device sent the FLOGI frame to the switch, and the switch probed the device.

To check for zoning problems

1. Enter the **cfgActvShow** command to determine if zoning is enabled.

   If zoning is enabled, it is possible that the problem is being caused by zoning enforcement (for example, two devices in different zones cannot see each other).
2. Confirm that the specific edge devices that need to communicate with each other are in the same zone.
   - If they are in the same zone, perform the following tasks:
     - Enter the **portCamShow** command on the host port to verify that the target is present.
     - Enter the **portCamShow** command on the target.
     - Enter the **nsZoneMember** command on the host and target to determine whether or not the Name Server is aware that these devices are zoned together
   - If they are not in the same zone and zoning is enabled, proceed to step 3.
3. Resolve zoning conflicts by putting the devices into the same zoning configuration.

   See "Correcting Zoning Setup Issues" on page 15-10 for additional information.

Restoring a Segmented Fabric

Fabric segmentation is generally caused by:

- Incompatible fabric parameters (see "To reconcile fabric parameters individually" on page 15-9).
- Incorrect PID setting (see "Configuring the PID Format" on page A-1).
- Incompatible zoning configuration (see "To check for zoning problems" on page 15-8).
- Domain ID conflict (see "To reconcile a domain ID conflict" on page 15-9).
- A switch in a secure fabric is not running Secure Fabric OS.

  See the *Secure Fabric OS Administrator's Guide* for additional information.

There are a number of settings that control the overall behavior and operation of the fabric. Some of these values, such as the domain ID, are assigned automatically by the fabric and can differ from one switch to another in the fabric. Other parameters, such as the BB credit, can be changed for specific applications or operating environments, but must be the same among all switches to allow the formation of a fabric.

The following fabric parameters must be identical for a fabric to merge:

- R_A_TOV
- E_D_TOV
- Data field size

- Sequence level switching
- Disable device probing
- Suppress class F traffic
- Per-frame route priority
- Long distance fabric (not necessary on Bloom-based fabrics)
- BB credit
- PID format

To reconcile fabric parameters individually

1. Log in to one of the segmented switches as admin (switch A).
2. Enter the **configShow** command.
3. Log in to another switch (switch B) in the same fabric as admin.
4. Enter the **configShow** command.
5. Compare the two switch configurations line by line and look for differences. Do this by comparing the two telnet windows or by printing the **configShow** output. Also verify that the fabric parameter settings (see page 15-8) are the same for *both* switches.
6. Connect to the segmented switch after the discrepancy is identified.
7. Disable the switch by entering the **switchDisable** command.
8. Enter the **configure** command to edit the fabric parameters for the segmented switch.

   See the *Fabric OS Command Reference* for more detailed information.
9. Enable the switch by entering the **switchEnable** command.

Alternatively, you can reconcile fabric parameters by entering the **configUpload** command for each switch.

To download a correct configuration

You can restore a segmented fabric by downloading a previously saved correct backup configuration to the switch. Downloading in this manner reconciles any discrepancy in the fabric parameters and allows the segmented switch to rejoin the main fabric. For details on uploading and downloading configurations, see Chapter 6, “Maintaining Configurations”.

To reconcile a domain ID conflict

If a domain ID conflict appears, the conflict is only reported at the point where the two fabrics are physically connected. However, there might be several conflicting domain IDs, which will appear as soon as the initial conflict is resolved.

Typically, the fabric automatically resolves domain conflicts during fabric merges or builds unless Insistent Domain ID (IDID) is configured. If IDID is enabled, switches that cannot be programmed with a unique domain ID are segmented out. Check each switch that has IDID configured and make sure their domain IDs are unique within the configuration.

Repeat this procedure until all domain ID conflicts are resolved:

1. Enter the **fabricShow** command on a switch from one of the fabrics.
2. In a separate telnet window, enter the **fabricShow** command on a switch from the second fabric.

3. Compare the **fabricShow** output from the two fabrics. Note the number of domain ID conflicts; there might be several duplicate domain IDs that will need to be changed. Determine which switches have domain overlap and change the domain IDs for each of those switches.
4. Chose the fabric on which to change the duplicate domain ID; connect to the conflicting switch in that fabric.
5. Enter the **switchDisable** command.
6. Enter the **switchEnable** command.

   This will enable the joining switch to obtain a new domain ID as part of the process of coming online. The fabric principal switch will allocate the next available domain ID to the new switch during this process.
7. Repeat step 4 through step 6 if additional switches have conflicting domain IDs.

Correcting Zoning Setup Issues

The types of zone configuration discrepancies that can cause segmentation are listed in Table 15-2.

Table 15-2 Types of Zone Discrepancies

Conflict Cause	Description
Configuration mismatch	Occurs when zoning is enabled in both fabrics and the zone configurations that are enabled are different in each fabric.
Type mismatch	Occurs when the name of a zone object in one fabric is also used for a different type of zone object in the other fabric. A zone object is any device in a zone.
Content mismatch	Occurs when the definition in one fabric is different from the definition of a zone object with the same name in the other fabric.

Table 15-3 summarizes commands that are useful for debugging zoning issues.

Table 15-3 Commands for Debugging Zoning

Command	Function
aliCreate	Use to create a zone alias.
aliDelete	Use to delete a zone alias.
cfgCreate	Use to create a zone configuration.
cfgShow	Displays zoning configuration.
defZone	Sets the default zone access mode to *No Access*, initializes a zoning transaction (if one is not already in progress), and creates the reserved zoning objects.
licenseShow	Displays current license keys and associated (licensed) products.
switchShow	Displays currently enabled configuration and any E_Port segmentations due to zone conflicts.

Table 15-3 Commands for Debugging Zoning (Continued)

Command	Function
zoneAdd	Use to add a member to an existing zone.
zoneCreate	Use to create a zone. Before a zone becomes active, the **cfgSave** and **cfgEnable** commands must be used.
zoneHelp	Displays help information for zone commands.
zoneShow	Displays zone information.

See "Administering Advanced Zoning" on page 20-1 for additional information about setting up zoning properly. Also see the *Fabric OS Command Reference* for details about zoning commands.

You can correct zone conflicts by using the **cfgClear** command to clear the zoning database.

Caution
This is a disruptive procedure.

To correct a fabric merge problem quickly

1. Determine which switch(es) have the incorrect zoning configuration; then, log in to the switches as admin.
2. Enter the **switchDisable** command.
3. Enter the **cfgDisable** command.
4. Enter the **cfgClear** command.

Caution
This command clears the zoning database on the affected switches.

5. Enter the **switchEnable** command. This forces a zone merge and populates the switches with the desired zoning database. The two fabrics will be merged again.

To correct a merge conflict without disrupting the fabric, first verify fabric merge problem, then edit zone configuration members, and then reorder the zone member list.

To verify a fabric merge problem

1. Enter the **switchShow** command to validate that the segmentation is due to a zone issue.
2. See Table 15-2 on page 15-10 to view the different types of zone discrepancies.

To edit zone configuration members

1. Log in to one of the switches in a segmented fabric as admin.
2. Enter the **cfgShow** command.
3. Print the output from the **cfgShow** command.
4. Start another telnet session and connect to the next fabric as an administrator.
5. Run the **cfgShow** command.
6. Print the output from the **cfgShow** command.

7. Compare the two fabric zone configurations line by line and look for an incompatible configuration.
8. Connect to one of the fabrics.
9. Run zone configure edit commands to edit the fabric zone configuration for the segmented switch (see Table 15-3 on page 15-10 for specific commands.

   If the zoneset members between two switches are not listed in the same order in both configurations, the configurations are considered a mismatch; this results in the switches being segmented in the fabric.

   For example:

   `[cfg1 = z1; z2]` is different from `[cfg1 = z2; z1]`, even though the members of the configuration are the same.

   One simple approach to making sure that the zoneset members are in the same order is to keep the members in alphabetical order.

To reorder the zone member list

1. Use the output from the **cfgShow** for both switches.
2. Compare the order that the zone members are listed. Members must be listed in the same order.
3. Rearrange zone members so that the configuration for both switches is the same. Arrange zone members in alphabetical order, if possible.

Recognizing MQ-WRITE Errors

An MQ error is a message queue error. Identify an MQ error message by looking for the two letters *M* and *Q* in the error message:

```
2004/08/24-10:04:42, [MQ-1004], 218,, ERROR, ras007, mqRead, queue = raslog-test-
string0123456-raslog, queue I
D = 1, type = 2
```

MQ errors can result in devices dropping from the SNS or can prevent a switch from joining the fabric. MQ errors are rare and difficult to troubleshoot; resolve them by working with the switch supplier.When MQ errors are encountered, issue the **supportSave** command to capture debug information about the switch; then, forward the **supportSave** data to the switch supplier for further investigation.

Correcting I^{2}C Bus Errors

I^{2}C bus errors generally indicate defective hardware or poorly seated devices or blades; the specific item is listed in the error message. See the *Fabric OS Message Reference* for information specific to the error that was received. Some CPT and Environmental Monitor (EM) messages contain I^{2}C-related information.

If the I^{2}C message does not indicate the specific hardware that might be failing, begin debugging the hardware, as this is the most likely cause. The next sections provide procedures for debugging the hardware.

To check fan components

1. Log in to the switch as user.
2. Enter the **fanShow** command.
3. Check the fan status and speed output.

If any of the fan speeds display abnormal RPMs, replace the fan FRU. You might first consider re-seating the fan FRU (unplug it and plug it back in).

To check the switch temperature

1. Log in to the switch as user.
2. Enter the **tempShow** command.
3. Check the temperature output.

   Look for indications of high or low temperatures.

To check the power supply

1. Log in to the switch as user.
2. Enter the **psShow** command.
3. Check the power supply status. Refer to the appropriate hardware reference manual for details regarding the power supply status.

   If any of the power supplies show a status other than OK, consider replacing the power supply as soon as possible.

To check the temperature, fan, and power supply

1. Log in to the switch as user.
2. Enter the **sensorShow** command. See the *Fabric OS Command Reference* for details regarding the sensor numbers.
3. Check the temperature output.

   Look for indications of high or low temperatures.
4. Check the fan speed output.

   If any of the fan speeds display abnormal RPMs, replace the fan FRU.
5. Check the power supply status.

If any of the power supplies show a status other than OK, consider replacing the power supply as soon as possible.

Correcting Device Login Issues

To try to pinpoint problems with device logins, use this procedure:

1. Log in to the switch as admin.
2. Enter the **switchShow** command; then, check for correct logins:

```
switch:admin> switchshow
switchName:     sw094135
switchType:     26.1
switchState:    Online
switchMode:     Native
switchRole:     Principal
switchDomain:   126
switchId:       fffc7e
switchWwn:      10:00:00:05:1e:34:00:69
zoning:         ON (cfg_em)
switchBeacon:   OFF
Port   Media Speed State
=========================
  0     id    N1   Online    E-Port  10:00:00:60:69:11:f9:fc "2800_116"
  1     id    1G   Online    E-Port  10:00:00:60:69:11:f9:fc "2800_116"
  2     id    N2   No_Light
  3     id    2G   No_Light
  4     id    N2   Online    E-Port  (Trunk port, master is Port   5)
  5     id    N2   Online    E-Port  10:00:00:05:1e:34:00:8b "Dazz125"
(downstream)(Trunk master)
  6     id    N2   No_Light
  7     id    N2   No_Light
  8     id    N1   Online    L-Port  4 public, 1 private, 1 phantom
  9     id    N2   No_Light
 10     id    N2   Online    G-Port
 11     id    N2   Online    F-Port  10:00:00:01:c9:28:c7:01
 12     id    N1   Online    L-Port  4 public, 1 private, 1 phantom
 13     --    N2   No_Module
 14     id    N2   Online    E-Port  (Trunk port, master is Port  15)
 15     id    N2   Online    E-Port  10:00:00:60:69:90:03:17 "TERM_113"
(downstream)(Trunk master)
```

3. Enter the **portCfgShow** command to see how the port is configured:

```
switch:admin> portcfgshow
Ports of Slot 8   16 17 18 19   20 21 22 23   24 25 26 27   28 29 30 31
-----------------+--+--+--+--+----+--+--+--+----+--+--+--+----+--+--+--
Speed             AN AN AN AN   AN AN AN AN   AN AN AN AN   AN AN AN AN
Trunk Port        ON ON ON ON   ON ON ON ON   .. ON ON ON   ON ON ON ..
Long Distance     .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
VC Link Init      .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Locked L_Port     .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Locked G_Port     .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Disabled E_Port   .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
ISL R_RDY Mode    .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
RSCN Suppressed   .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Persistent DisableON ON ON ON   ON ON ON ON   .. ON ON ON   ON ON ON ON
NPIV capability   ON ON ON ON   ON ON ON ON   ON ON ON ON   ON ON ON ON
EX Port           .. .. .. ..   .. .. .. ..   ON .. .. ..   .. .. .. ..

                  where AN:AutoNegotiate, ..:OFF, ??:INVALID.
                        LM:L0.5
```

4. Enter the **portErrShow** command; then, check for errors that can cause login problems:

```
switch:admin> porterrshow
frames  enc  crc  too  too  bad  enc disc link loss loss frjt fbsy
       tx   rx   in  err shrt long  eof  out   c3 fail sync  sig
     ========================================================================
 64:    0    0    0    0    0    0    0    0    0    0    0    1    0    0
 65:    0    0    0    0    0    0    0    0    0    0    0    1    0    0
 66:    0    0    0    0    0    0    0    0    0    0    0    1    0    0
 67:    0    0    0    0    0    0    0    0    0    0    0    1    0    0
 68:    0    0    0    0    0    0    0    0    0    0    0    1    0    0
 69:    0    0    0    0    0    0    0    0    0    0    0    1    0    0
 70:    0    0    0    0    0    0    0    0    0    0    0    1    0    0
 71:    0    0    0    0    0    0    0    0    0    0    0    1    0    0
 72:    0    0    0    0    0    0    0    0    0    0    0    1    0    0
 73:    0    0    0    0    0    0    0    0    0    0    0    1    0    0
 74:    0    0    0    0    0    0    0    0    0    0    0    1    0    0
 75:    0    0    0    0    0    0    0    0    0    0    0    1    0    0
 76:    0    0    0    0    0    0    0    0    0    0    0    1    0    0
 77:    0    0    0    0    0    0    0    0    0    0    0    1    0    0
 78:    0    0    0    0    0    0    0    0    0    0    0    1    0    0
 79:    0    0    0    0    0    0    0    0    0    0    0    1    0    0
 80:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
 81:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
 82:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
 83:   58k 111k   0    0    0    0    0    0  231    0    0    0    0    0
 84:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
 85:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
 86:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
 87:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
 88:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
 89:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
 90:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
 91:  3.5g 3.5g   2    5    0    5    0    6  2.0k   5    5    5    0    0
 92:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
 93:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
```

- A high number of errors relative to the frames transmitted and frames received can indicate a marginal link (see “Correcting Marginal Links” on page 15-24 for additional information).
- A steadily increasing number of errors can indicate a problem. Track errors by sampling the port errors every five or ten minutes.

5. Enter the **portFlagsShow** command; then, check to see how a port has logged in and where a login failed (if a failure occurred):

```
switch:admin> portflagsshow
Slot Port SNMP      Physical   Flags
-----------------------------------
   7    0 Offline   No_Module  PRESENT U_PORT LED
   7    1 Offline   No_Module  PRESENT U_PORT LED
   7    2 Offline   No_Module  PRESENT U_PORT LED
   7    3 Offline   No_Module  PRESENT U_PORT LED
   7    4 Offline   No_Module  PRESENT U_PORT LED
   7    5 Offline   No_Module  PRESENT U_PORT LED
   7    6 Offline   No_Module  PRESENT U_PORT LED
   7    7 Offline   No_Module  PRESENT U_PORT LED
   7    8 Offline   No_Module  PRESENT U_PORT LED
   7    9 Offline   No_Module  PRESENT U_PORT LED
   7   10 Offline   No_Module  PRESENT U_PORT LED
   7   11 Offline   No_Module  PRESENT U_PORT LED
   7   12 Offline   No_Module  PRESENT U_PORT LED
   7   13 Offline   No_Module  PRESENT U_PORT LED
   7   14 Offline   No_Module  PRESENT U_PORT LED
   7   15 Offline   No_Module  PRESENT U_PORT LED
   7   16 Offline   No_Module  PRESENT U_PORT LED
   7   17 Offline   No_Module  PRESENT U_PORT LED
   7   18 Offline   No_Module  PRESENT U_PORT LED
   7   19 Offline   No_Module  PRESENT U_PORT LED
   7   20 Offline   No_Module  PRESENT U_PORT LED
   7   21 Offline   No_Module  PRESENT U_PORT LED
   7   22 Offline   No_Module  PRESENT U_PORT LED
   7   23 Offline   No_Module  PRESENT U_PORT LED
   7   24 Offline   No_Module  PRESENT U_PORT LED
   7   25 Offline   No_Module  PRESENT U_PORT LED
   7   26 Offline   No_Module  PRESENT U_PORT LED
   7   27 Offline   No_Module  PRESENT U_PORT DISABLED LED
   7   28 Offline   No_Module  PRESENT U_PORT LED
   7   29 Offline   No_Module  PRESENT U_PORT LED
   7   30 Offline   No_Module  PRESENT U_PORT LED
   7   31 Offline   No_Module  PRESENT U_PORT LED
   8    0 Offline   No_Module  PRESENT U_PORT DISABLED LED
   8    1 Offline   No_Module  PRESENT U_PORT DISABLED LED
   8    2 Offline   No_Module  PRESENT U_PORT DISABLED LED
   8    3 Online    In_Sync    PRESENT ACTIVE E_PORT G_PORT U_PORT LOGICAL_ONLINE
LOGIN LED
   8    4 Offline   No_Module  PRESENT U_PORT DISABLED LED
   8    5 Offline   No_Module  PRESENT U_PORT DISABLED LED
   8    6 Offline   No_Module  PRESENT U_PORT DISABLED LED
   8    7 Offline   No_Module  PRESENT U_PORT DISABLED LED
   8    8 Offline   No_Module  PRESENT U_PORT DISABLED LED
   8    9 Offline   No_Module  PRESENT U_PORT DISABLED LED
   8   10 Offline   No_Module  PRESENT U_PORT DISABLED LED
   8   11 Online    In_Sync    PRESENT ACTIVE E_PORT G_PORT U_PORT LE_PORT EX_PORT
LOGICAL_ONLINE LOGIN LED
   8   12 Offline   No_Module  PRESENT U_PORT DISABLED LED
   8   13 Offline   No_Module  PRESENT U_PORT DISABLED LED
   8   14 Offline   No_Module  PRESENT U_PORT DISABLED LED
   8   15 Online    In_Sync    PRESENT ACTIVE F_PORT G_PORT U_PORT LOGICAL_ONLINE
LOGIN NOELP ACCEPT
```

6. Enter the **portLogDumpPort** *portid* command where the port ID is the port number; then, view the device to switch communication.

```
switch:admin> portlogdumpport 10
time          task       event  port cmd  args
-----------------------------------------------
12:38:21.590  SPEE       sn      10   WS  00000000,00000000,00000000
12:38:21.591  SPEE       sn      10   WS  000000ee,00000000,00000000
12:38:21.611  SPEE       sn      10   WS  00000001,00000000,00000000
12:38:21.871  SPEE       sn      10   NC  00000002,00000000,00000001
12:38:21.872  LOOP       loopscn 10  LIP  8002
12:38:22.171  LOOP       loopscn 10  TMO  2
12:38:22.171  INTR       pstate  10  LF2
12:38:22.172  INTR       pstate  10  OL2
12:38:22.172  INTR       pstate  10  LR3
12:38:22.172  INTR       pstate  10   AC
12:38:22.172  PORT       scn     10   11  00000000,00000000,00000002
12:38:22.311  PORT       scn     10    1  00000000,00000000,00000001
12:38:22.311  PORT       debug   10       00000001,00654320,00000001,00000000
12:38:22.311  PORT       debug   10       00000001,00654320,00000002,00000000
12:38:22.311  PORT       debug   10       00000001,00654320,00000003,00000000
12:38:22.313  PORT       Tx      10  164  02fffffd,00fffffd,025effff,10000000
12:38:22.314  PORT      debug   10      00000001,00654320,00000003,00000000   *
7
12:38:28.312  PORT       Tx      10  164  02fffffd,00fffffd,028fffff,10000000
12:38:34.312  PORT       Tx      10  164  02fffffd,00fffffd,0293ffff,10000000
12:38:40.312  PORT       Tx      10  164  02fffffd,00fffffd,0299ffff,10000000
12:38:46.312  PORT       Tx      10  164  02fffffd,00fffffd,029bffff,10000000
12:38:52.312  PORT       Tx      10  164  02fffffd,00fffffd,029dffff,10000000
12:38:58.312  PORT       Tx      10  164  02fffffd,00fffffd,02acffff,10000000
12:39:04.322  INTR       pstate  10  LR1
12:39:04.323  INTR       pstate  10  LR3
12:39:04.323  INTR       pstate  10   AC
12:39:04.323  PORT       scn     10   11  00000000,00000000,00000002
switch:admin>>
```

Note

See “Viewing the Port Log” on page 14-11 for overview information about a **portLogDump**.

Identifying Media-Related Issues

This section provides procedures that help pinpoint any media-related issues in the fabric. The tests listed in Table 15-4 are a combination of *structural* and *functional* tests that can be used to provide an overview of the hardware components and help identify media-related issues.

- *Structural* tests perform basic testing of the switch circuit. If a structural test fails, replace the main board or port blade.
- *Functional* tests verify the intended operational behavior of the switch by virtue of running frames through ports or bypass circuitry.

Table 15-4 Component Test Descriptions

Test Name	Operands	Checks
crossporttest	[**-nframes** *count*] [**-lb_mode** *mode*][**-spd_mode** *mode*] [**-gbic_mode** *mode*] [**-norestore** *mode*] [**-ports** *itemlist*]	Functional test of port external transmit and receive path. The **crossport** is set to loopback using an external cable by default. However, this command can be used to check internal components by setting the *lb* operand to 5.
fporttest	[**-nframes** *count*] **[-ports** *itemlist*] [**-seed** *payload_pattern*] [**-width** *pattern_width*] [**-size** *pattern_size*]	Tests component to/from and HBA. Used to test online F_Port devices, N_Port devices, SFPs, and GBICs.
loopporttest	[**-nframes** *count*] [**-ports** *itemlist*][**-seed** *payload_pattern*] [**-width** *pattern_width*]	Only tests components attached to a switch that are on a FC-AL.
spinfab	*[* **-nmegs** *count] [* **-ports** *itemlist] [***-setfail** *mode]*	Tests switch-to-switch ISL cabling and trunk group operations.

The following procedures are for checking switch-specific components.

To test a port's external transmit and receive path

1. Connect to the switch and log in as admin.
2. Connect the port you want to test to any other switch port with the cable you want to test.
3. Enter the **crossPortTest** command with the following operands as shown in this example (this is a partial list; see the *Fabric OS Command Reference* for additional command information):

   [**-nframes** *count*]—Specify the number of frames to send.

   [**-lb_mode** *mode*]—Select the loopback point for the test.

   [**-spd_mode** *mode*]—Select the speed mode for the test.

   [**-ports** *itemlist*]—Specify a list of user ports to test.

```
switch:admin> crossporttest
Running Cross Port Test .... passed.
```

To test a switch's internal components

1. Connect to the switch and log in as admin.
2. Connect the port you want to test to any other switch port with the cable you want to test.
3. Enter the **crossporttest -lb_mode 5** command where *5* is the operand that causes the test to be run on the internal switch components (this is a partial list—see the *Fabric OS Command Reference* for additional command information):

   [**-nframes** *count*]—Specify the number of frames to send.

   [**-lb_mode** *mode*]—Select the loopback point for the test.

   [**-spd_mode** *mode*]—Select the speed mode for the test.

   [**-ports** *itemlist*]—Specify a list of user ports to test.

To test components to and from the HBA

1. Connect to the switch and log in as admin.
2. Enter the **fPortTest** command as shown in this example (see the *Fabric OS Command Reference* for information on the command options).

```
switchname:admin> fporttest 100,8,0xaa55,2, 512
 Will use pattern: aa55  aa55  aa55 aa55  aa55  aa55 ...
 Running fPortTest ........
 port 8 test passed.
 value = 0
```

   This example executed the **fPortTest** command 100 times on port 8 with payload pattern 0xaa55, pattern width 2 (meaning word width) and a default payload size of 512 bytes.

See Table 15-5 for a list of additional tests that can be used to determine the switch components that are not functioning properly. See the *Fabric OS Command Reference* for additional command information.

Table 15-5 Switch Component Tests

Test	Function
portloopbacktest	Functional test of port N to N path.
portregtest	A read and write test of the ASIC SRAMs and registers.
spinsilk	Functional test of internal and external transmit and receive paths at full speed.
sramretentiontest	Verifies that the data written into the miscellaneous SRAMs in the ASIC are retained after a 10-second wait.
crossporttest	Verifies the functional components of the switch.
turboramtest	Verifies that the on chip SRAM located in the 2 Gbit/sec ASIC is using the Turbo-Ram BIST circuitry. These same SRAMs are tested by **portregtest** and **sramretentiontest** using PCI operations, but for this test the BIST controller is able to perform the SRAM write and read operations at a much faster rate.
statstest	Verifies that the ASIC statistics counter logic.
Related Switch Test Option:	
itemlist	Restricts the items to be tested to a smaller set of parameter values that you pass to the switch.

Correcting Link Failures

A link failure occurs when a server or storage is connected to a switch, but the link between the server/storage and the switch does not come up. This prevents the server/storage from communicating through the switch.

If the **switchShow** command or LEDs indicate that the link has not come up properly, use one or more of the following procedures.

The port negotiates the link speed with the opposite side. The negotiation usually completes in 1-2 seconds; however, sometimes the speed negotiation fails.

Note

Skip this procedure if the port speed is set to a static speed through the **portCfgSpeed** command.

To determine if the negotiation was successfully completed

1. Enter the **portCfgShow** command to display the port speed settings of all the ports.
2. Enter the **switchShow** command to determine if the port has module light.
3. Determine whether or not the port at 1Gig/sec completes by entering the **portCfgSpeed** command. Then change the port speed to 2 Gig/sec. This should correct the negotiation by setting to one speed.
4. Enter the **portLogShow** or **portLogDump** command.
5. Check the events area of the output. The first example is 1 Gbit/sec and the second example is 2 Gbit/sec:

```
14:38:51.976  SPEE sn <Port#>   NC  00000001,00000000,00000001
```

```
14:39:39.227  SPEE       sn      <Port#>   NC  00000002,00000000,00000001
```

   - *sn* indicates a speed negotiation.
   - *NC* indicates negotiation complete.
   - *01* or *02* indicate the speed that has been negotiated.

   If these fields do not appear, proceed to the step 6.

6. Correct the negotiation by entering the **portCfgSpeed** [*slotnumber*/]*portnumber*, *speed_level* command if the fields in step 5 do not appear.

To check for a loop initialization failure

1. Verify the port is an L_Port.
   a. Enter the **switchShow** command.
   b. Check the comment field of the output to verify that the switch port indicates an L_Port. If a loop device is connected to the switch, the switch port must be initialized as an L_Port.
   c. Check to ensure that the port state is online; otherwise, check for link failures.
2. Verify the loop initialization *if* the port is not an L_port.
   a. Enter the **portLogShow** or **portLogDump** command.
   b. Check argument number four for the *LISA frame* (0x11050100).

```
switch:admin> portlogdumpport 4
time          task       event  port cmd  args
-----------------------------------------------
11:40:02.078  PORT       Rx3     23   20  22000000,00000000,ffffffff,11050100
Received LISA frame
```

   The *LISA frame* indicates that the loop initialization is complete.
3. Skip point-to-point initialization.

   The switch changes to point-to-point initialization after the Loop Initialization Soft Assigned (LISA) phase of the loop initialization. This behavior sometimes causes trouble with old HBAs. If this is the case, then:

   Skip point-to-point initialization by using the **portCfgLport** Command.

To check for a point-to-point initialization failure:

1. Enter the **switchShow** command to confirm that the port is active and has a module that is synchronized.

   If a fabric device or another switch is connected to the switch, the switch port must be online.
2. Enter the **portLogShow** or **portLogDump** commands.
3. Verify that the event area for the port state entry is *pstate*. The command entry *AC* indicates that the port has completed point-to-point initialization.

```
switch:admin> portlogdumpport 4
time          task       event  port cmd  args
-----------------------------------------------
11:38:21.726  INTR       pstate   4   AC
```

4. Skip over the loop initialization phase.

   After becoming an active port, the port becomes an F_Port or an E_Port depending on the device on the opposite side. If the opposite device is a fabric device, the port becomes an F_Port. If the opposite device is another switch, the port becomes an E_Port.

   Some fabric devices have problems with loop initialization. If this is evident, enter the **portCfgLPort** *port #***, 1** command; this will force the port to try and come up as a public L-Port.

To correct a port that has come up in the wrong mode

1. Enter the **switchShow** command.
2. Refer to the comment fields (see Table 15-6) and follow the suggested actions.

Table 15-6 SwitchShow Output and Suggested Action

Output	Suggested Action
Disabled	Check the output from the **switchShow** command to determine whether or not the switch is disabled. If the port is disabled (for example, due to persistent disable or security reasons), attempt to resolve the issue and then enter the **portEnable** command.
Bypassed	Check the output from the **switchShow** command to determine whether or not the port is testing.
Loopback	Check the output from the **switchShow** command to determine whether or not the port is testing.
E_Port	If the opposite side is not another switch, the link has come up in a wrong mode. Check the output from the **portLogShow** or **PortLogDump** commands and identify the link initialization stage where the initialization procedure went wrong.
F_Port	If the opposite side of the link is a fabric device, the link has come up in a wrong mode. Check the output from **portLogShow** or **PortLogDump** commands.
G_Port	The port has not come up as an E_Port or F_Port. Check the output from **portLogShow** or **PortLogDump** commands and identify the link initialization stage where the initialization procedure went wrong.
L_Port	If the opposite side is *not* a loop device, the link has come up in a wrong mode. Check the output from **portLogShow** or **PortLogDump** commands and identify the link initialization stage where the initialization procedure went wrong.

Correcting Marginal Links

A marginal link involves the connection between the switch and the edge device. Isolating the exact cause of a marginal link involves analyzing and testing many of the components that make up the link (including the switch port, switch SFP, cable, the edge device, and the edge device SFP).

To troubleshoot a marginal link

1. Enter the **portErrShow** command.

```
switch:admin> porterrshow
       frames  enc  crc  too  too  bad  enc disc link loss loss frjt fbsy
      tx   rx   in  err shrt long  eof  out   c3 fail sync sig
sig=================================================================
 0:   22   24    0    0    0    0    0  1.5m   0    7    3    0    0    0
 1:   22   24    0    0    0    0    0  1.2m   0    7    3    0    0    0
 2:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
 3:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
 4:  149m  99m   0    0    0    0    0  448    0    7    6    0    0    0
 5:  149m  99m   0    0    0    0    0  395    0    7    6    0    0    0
 6:  147m  99m   0    0    0    0    0  706    0    7    6    0    0    0
 7:  150m  99m   0    0    0    0    0  160    0    7    5    0    0    0
 8:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
 9:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
10:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
11:    0    0    0    0    0    0    0    0    0    0    0    2    0    0
12:    0    0    0    0    0    0    0    0    0    0    0    2    0    0
13:    0    0    0    0    0    0    0    0    0    0    0    2    0    0
14:    0    0    0    0    0    0    0    0    0    0    0    2    0    0
15:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
32:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
33:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
34:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
35:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
36:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
37:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
38:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
39:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
40:   99m 146m   0    0    0    0    0  666    0    6  796    7    0    0
41:   99m 149m   0    0    0    0    0   15k   0    2  303    4    0    0
42:   99m 152m   0    0    0    0    0  665    0    2  221    5    0    0
43:   99m 147m   0    0    0    0    0   16k   0    2  144    4    0    0
44:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
45:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
46:    0    0    0    0    0    0    0    0    0    0    0    2    0    0
47:    0    0    0    0    0    0    0    0    0    0    0    0    0    0
switch:admin>
```

2. Establish if there are a relatively high number of errors (such as CRC errors or ENC_OUT errors), or if there are a steadily increasing number of errors to confirm a marginal link.
3. If you suspect a marginal link, isolate the areas by moving the suspected marginal port cable to a different port on the switch. Reseating of SFPs may also cure marginal port problems.

   If the problem stops or goes away, the switch port or the SFP is marginal (proceed to step 4)

   If the problem does *not* stop or go away, see step 7.

4. Replace the SFP on the marginal port.
5. You will need an adapter to run the loopback test for the SFP. Otherwise, run the **portloopbacktest** on the marginal port using the loopback mode *lb=5*. See the *Fabric OS Command Reference* for additional information.

Table 15-7 Loopback Modes

Loopback Mode	Description
1	Port Loopback (loopback plugs)
2	External (SERDES) loopback
5	Internal (parallel) loopback (indicates no external equipment)
7	Back-end bypass & port loopback
8	Back-end bypass & SERDES loopback
9	Back-end bypass & internal loopback

6. Check the results of the loopback test and proceed as follows:
   - If the loopback test failed, the port is bad. Replace the port blade.
   - If the loopback test did not fail, the SFP was bad.
7. Optionally, to rule out cabling issues:
   a. Insert a new cable in to the suspected marginal port.
   b. Enter the **portErrShow** command to determine if a problem still exists.
      - If the **portErrShow** output displays a normal number of generated errors, the issue is solved.
      - If the **portErrShow** output still displays a high number of generated errors, follow the troubleshooting procedures for the Host or Storage device.

Inaccurate Information in the System Message Log

In rare instances, events gathered by the *track change* feature can report inaccurate information to the system message log.

For example, a user enters a correct user name and password, but the login was rejected because the maximum number of users had been reached. However, when looking at the system message log, the login was reported as successful.

If the maximum number of switch users has been reached, the switch will still perform correctly in that it will reject the login of additional users (even if they enter correct user name and password information).

However, in this limited example, the Track Change feature will report this event inaccurately to the system message log; it will appear that the login was successful. This scenario only occurs when the maximum number of users has been reached; otherwise, the login information displayed in the system message log should reflect reality.

See "Tracking and Controlling Switch Changes" for information regarding enabling and disabling track changes (TC).

Recognizing the Port Initialization and FCP Auto Discovery Process

The steps in the port initialization process represent a protocol used to discover the type of connected device and establish the port type. The possible port types are as follows:

- U_Port—Universal FC port. This port type is the base Fibre Channel port type and all unidentified, or uninitiated ports are listed as U_Ports.
- FL_Port—Fabric Loop port. This port connects both public and private loop devices.
- G_Port—Generic port. This port acts a transition port for non-loop fabric capable devices (E_Port / F_Port).
- E_Port—Expansion port. This port type is assigned to ISL links.
- F_Port—Fabric port. This port is assigned to fabric capable devices.
- EX_Port—A type of E_Port that connects a Fibre Channel router to an edge fabric. From the point of view of a switch in an edge fabric, an EX_Port appears as a normal E_Port. It follows applicable Fibre Channel standards as other E_Ports. However, the router terminates EX_Ports rather than allowing different fabrics to merge as would happen on a switch with regular E_Ports.
- VE_Port—Functions somewhat like an E_Port, but terminates at the switch and does not propagate fabric services or routing topology information from one edge fabric to another.
- VEX_Port—A type of VE_Port that connects a Fibre Channel router to an edge fabric. From the point of view of a switch in an edge fabric, an VEX_Port appears as a normal VE_Port. It follows the same Fibre Channel protocol as other VE_Ports. However, the router terminates VEX_Ports rather than allowing different fabrics to merge as would happen on a switch with regular VE_Ports.

The Brocade FCP auto discovery process enables private storage devices that accept PRLI to communicate in a fabric.

If device probing is enabled, the embedded port PLOGIs and attempts a PRLI into the device to retrieve information to enter into the Name Server. This enables private devices that do not FLOGI but accept PRLI to be entered in the Name Server and receive full fabric citizenship. Private devices that accept PRLI represent a majority of storage targets. Private hosts require the QuickLoop feature, which is not available in Fabric OS v4.0.0 or later.

A fabric-capable device will implicitly register information with Name Server during a FLOGI. These devices will typically register information with the Name Server before querying for a device list. The embedded port will still PLOGI and attempt PRLI with these devices.

You can view the Name Server table in Web Tools by clicking **Name Server** in the fabric toolbar. See the *Web Tools Administrator's Guide* for more information.

Port Mirroring

Port mirroring lets you configure a switch port as an analyzer port to mirror a specific source port and destination port traffic passing though any switch port. This is a useful way to troubleshoot without bringing down the host and destination links to insert an inline analyzer.

Port mirroring captures traffic between two devices. It mirrors only the frames containing the SID/DID to the mirror port. Because of the way it handles mirroring, a single mirror port can mirror multiple mirror connections. This also means that the port cannot exceed the maximum bandwidth of the mirror port. Attempts to mirror more traffic than available bandwidth result in the port mirror throttling the SID/DID traffic so that traffic does not exceed the maximum available bandwidth.

Port mirroring is supported between VE_ports (VE_Port to VE_Port) with FCIP and no routing, and over VEx _Ports and Ex_Ports. The mirror port can be any port located on the same switch as the source identifier (SID).

Use port mirroring to detect missing frames, which might occur with zoning issues or hold timeouts, capture protocol errors, and capture ULP traffic (SCSI/FICON). This feature cannot be used on embedded switch traffic.

Port mirroring is only available using the FOS 5.2.0 CLI and is not available through Web Tools. For a complete list of port mirroring commands, see the *Fabric OS Command Reference*.

In order to ensure proper failover in HA configurations, both the active and the standby control processors (CP) must have firmware version 5.2 installed and running. If the OS on the standby CP does not support mirroring, failing over the standby CP could cause the HA failover to fail.

Supported Hardware

Port mirroring is supported on Condor-based ASIC platforms, including:

- SilkWorm 4100
- SilkWorm 4900
- SilkWorm 7500
- SilkWorm 48000 with chassis option 5

Port mirroring can be used on the following blades within a chassis:

- FC4-32 32-port blade
- FC4-16 16-port blade
- *FC4-48 48-port blade
- FR4-18i routing & FCIP blade

- FC4-16IP iSCSI blade on FC ports only

  The FC4-48 implements port pairing, meaning that two ports share the same area. Port pairing uses a single area to map to two physical ports. A frame destined to the secondary port is routed to the primary port. The primary port's filtering zone engine is used to redirect the frame to the secondary port. Port mirroring uses the port filter zone engine to redirect the frames to the mirror port. If two F_Ports share the same area, both ports cannot be part of a mirror connection. One of the two ports can be part of the connection as long as the other port is offline. Supported port configurations are shown in Table 15-8.

  Table 15-8 Port Combinations for Port Mirroring

Primary Port	Secondary Port	Supported
F_Port	F_Port	No
F_Port	Offline	Yes
Offline	F_Port	Yes
F_Port	E_Port	Yes
E_Port	F_Port	Yes
E_Port	E_Port	No

  If IOD is enabled, adding or deleting a port mirror connection causes a frame drop. Port mirroring reroutes a given connection to the mirror port, where the mirror traffic takes an extra route to the mirror port. When the extra route is removed, the frames between the two ports goes directly to the destination port. Since the frames at the mirror port could be queued at the destination port behind those frames that went directly to the destination port, port mirroring drops those frames from the mirror port when a connection is disabled. If IOD has been disabled, port mirroring does not drop any frames but displays an IOD error.

- A port cannot be mirrored to multiple locations. If you define multiple mirror connections for the same F_Port, all the connections must share the same mirror port.
- Local switches cannot be mirrored because FICON CUP frames to a local switch are treated as well-known addresses or embedded frame traffic.
- Using firmware download to downgrade to previous Fabric OS releases that do not support port mirroring requires that you remove all port mirroring connections. If you downgrade to a previous versions of Fabric OS, you cannot proceed until the mirroring connections are removed.

How Port Mirroring Works

Port mirroring reroutes the data frames between two devices to the mirror port. Rerouting introduces latency for the data flow. The latency depends on the location of the mirror port.

For a given port, the traffic received from the point of view of the switch can be captured before leaving this ASIC. Each user port is connected to an ASIC port. The user port's ingress traffic is routed to another user port on this chip, uplinks to the core switch, or E_Ports to remote domains. Mirroring is used on user ports, uplink ports (C_Ports), or E_Ports to capture the received traffic.

The user port's destination to source traffic can come from any of the E_Ports or user ports, either on the the same chip or a different chip.

When the destination port is on a different CP or is a remote switch port, port mirroring installs ingress (traffic received) filters on the C_Ports. If the destination port is on the same ASIC, an egress (traffic sent) filter is installed on the source port.

A port configured as a mirror port is configured so that the transmitter is looped back to the receiver. In addition, the transmit traffic is sent out of the transmit port. You could connect an analyzer to this port to collect all the traffic exiting this port. The traffic looped back will then be routed to the final destination. Traffic received externally from this port does not affect loopback traffic.

The performance server daemon (psd) maintains the port mirroring database, which contains the source identifier, destination identifier, and a mirror port. Adding a new mirror connection adds the connection to the database, where the information is saved. Deleting the connection removes the mirror connection from the database.

F_Port Connection—A port mirror connection containing both the source identifier and destination identifier on the same chip is called a “local chip connection.” In this type of connection, only two ports are involved to capture the sent and received traffic. The destination port mirrors the received (from the switch’s point of view) traffic.

- Traffic is received at the source port and the switch routes these frames to the destination port.
- The destination port has a port mirror, which redirects matching frames to the mirror port.
- The mirror port then routes those frames it receives back to the destination port.

FC_Port Connection—A port mirror connection between two F_Ports on different CPs but in the same switch, as in core-edge platforms. If one port is sending traffic to another port on a different chip, the port must send the traffic to the core to reach the edge ports located on another chip.

- Traffic is received at the source port destined to another port on a different chip. The frame is routed to one of the C_Ports connected to the core.
- All C_Ports have a port mirror filter to redirect those frames that match the filter to the mirror port.
- 3. The mirror port then routes those frames it receives back to a C_Port.

FCE_Port Connection—Port mirror connections on which the source identifier on the local switch and destination identifier is on a different switch is called an FCE_Port connection. In this type of connection, the source port, all the E-ports located on the same chip as the source port and all C_Ports located on the source chip are involved with this connection type. In addition, the C_Ports on other blades that have E_Ports are involved in this type of connection too.

The received traffic can be sent to any E_Port, regardless of whether the port is on the same or another chip. The port mirrors the received traffic on all C_Ports and all E_Ports on that source identifier's chip.

- Traffic is received at the source port destined to a remote port. The switch routes these frames to either a local E+Ports on this chip or to the C_Ports to reach another E-ports on a different blade.
- The C_Ports and E_ports have a port mirror filter on these ports to redirect those frames that match the filter to the mirror port.
- The mirror port then routes those frames it receives back to the either the C_ports or E_ports port.

There are two types of transmit filter installation

- If the E_Port is on the same chip, port mirroring installs an egress (transmitted information) filter on the source port.
- If the E_Port is on a different chip, port mirror installs the filter on the C_Ports of the other chip.

To better explain how the transmit filter works on each of these types, the method used for both types is described as follows:

- Traffic is received at the E_Ports destined to a source port. The switch routes these frames to the source port (if the E_Port is on the same chip) or to the C_Ports destined to the source port (if the E_Ports are on different chips).
- On the same chip, the source port has a mirror to redirect any matching frames to the mirror port. On different chips, the C_Ports redirect any matching frames to the mirror port.
- The mirror port routes those frames it receives back to the source port (on the same chip) or to the C_Ports (on a different chip).

When a port goes offline, the mirror connections that are related to this port will be removed. When the port goes online, if all ports of the connection are online, the mirror connections that are related to this port will be restored.

For remote ports, connections that are mirroring traffic from a remote domain, when the remote domain becomes unreachable, the mirror connections related to this domain are removed. When the remote domain becomes reachable, the mirror connection will only be added if all the ports of the connection are online too.

Port Mirroring Considerations

Before creating port mirror connections, consider the following limitations:

- A mirror port can be any port on the same switch as the source identifier port.
- Only one domain can be mirrored per chip; after a domain is defined, only mirror ports on the defined domain can be used.

  For example, in a three-domain fabric containing switches 4100A, 4100B, and 4100C, a mirror connection that is created between 4100A and 4100B only allows 4100A to add mirror connections for those ports on 4100B. To mirror traffic between 4100A and 4100C, add a mirror connection on 4100C. The first connection defines the restriction on the domain, which can be either the local domain or a remote domain.
- A switch that is capable of port mirroring can support a maximum of four mirror connections.

  Each FDB defines an offset to search. Each offset can have up to four values that can be defined for a filter. If any of the four values match, the filter will match.
- Mirror port bandwidth limits mirror connections.

  The bandwidth of the mirror port is unidirectional. The host (SID) talks to multiple storage devices (DIDs) and does not send full line rate to a single target. A mirror port configured at 2GB can only support up to 2GB of traffic. A normal 2G F_Port is bidirectional and can support up to 4GB of traffic (two to transmit and two to receive). If the mirror port bandwidth is exceeded, the receiver port is not returned any credits and the devices in the mirror connection see degraded performance.
- Deleting a port mirroring connection with In Order Deliver (IOD) enabled causes frame drop between two endpoints.
- Using the firmware download procedure to downgrade to previous Fabric OS releases that do not support port mirroring requires that you remove all the port mirroring connections. If you downgrade to a previous versions of Fabric OS, you cannot proceed until the mirroring connections are removed.

Creating, Deleting, and Displaying Port Mirroring

The following section describes how to use the port mirroring feature in the fabric.

The method for adding a port mirror connection between two local switch ports and between a local switch port and a remote switch port is the same.

To add a port mirror connection

1. Log in to the switch as admin.
2. Type **portMirror --add** *slotnumber/portnumber SourceID DestID*

The configuration database keeps information about the number of port mirror connections configured on a switch, the number of chunks of port mirroring data that are stored, and the chunk number. When removing a mirror connection, always use this method to ensure that the data is cleared. Deleting a connection removes the information from the database.

To delete a port mirror connection between two local switch ports or a local and a remote switch port

1. Log in to the switch as admin.
2. Type **portMirror --del** *SourceID DestID***.**

   For example, to delete the port mirror connection on mirror port 2, you might type:

   ```
   portMirror --del 0x011400 0x240400
   ```

To display port mirror connections

1. Log in to the switch as admin.
2. Type **portMirror --show**

   You should see output similar to the following:

```
switch:admin> portmirror --show

Number of mirror connection(s) configured: 4

Mirror_Port  SID       DID       State
----------------------------------------
18           0x070400  0x0718e2  Enabled
18           0x070400  0x0718e3  Enabled
18           0x070400  0x0718ef  Enabled
18           0x070400  0x0718e0  Enabled
```

The **switchShow** command output shows the mirror port as shown in the following example.

```
switch:admin> switchshow
switchName:ESS118
switchType: 42.2
switchState: Online
switchMode: Native
switchRole:Subordinate
switchDomain:   121 switchId:fffc79
switchWwn:10:00:00:60:69:e4:00:a0
zoning:ON (c)
switchBeacon:   OFF blade2 Beacon:  OFF

Area Slot Port Media Speed State
====================================
16  2 0   --N4   No_Module
17  2 1   idN2   No_Light
18  2 2   idN2   No_Light
19  2 3   idN2   No_Light
20  2 4   --N4   No_Module
21  2 5   idN2   No_Light
22  2 6   idN2   No_Light
23  2 7   idN2   No_Light
24  2 8   idN1   OnlineL_Port
(output truncated)
156 2 28   --N4   No_Module
157 2 29   idN2   No_Light
158 2 30   --N4   No_Module
159 2 31   idN2   OnlineMirror Port
```

Administering NPIV

Chapter 16

This chapter describes the concepts and procedures for administering N-Port ID Virtualization (NPIV) and contains the following topics:

- "About NPIV" on page 16-1
- "Enabling and Disabling NPIV" on page 16-2
- "Configuring NPIV" on page 16-2
- "Viewing NPIV Port Configuration Information" on page 16-3

About NPIV

N-Port ID Virtualization (NPIV) enables a single Fibre Channel protocol port to appear as multiple, distinct ports, providing separate port identification within the fabric for each operating system image behind the port (as if each operating system image had its own unique physical port). NPIV assigns a different virtual port ID to each Fibre Channel protocol device. NPIV is designed to enable you to allocate virtual addresses without impacting your existing hardware implementation. The virtual port has the same properties as an N_Port, and is therefore capable of registering with all services of the fabric.

Each NPIV device has a unique device PID, Port WWN, and Node WWN, and should act the same as all other physical devices in the fabric; in other words, multiple virtual devices emulated by NPIV appear no different than regular devices connected to a non-NPIV port. The same zoning rules apply to NPIV devices as non-NPIV devices. Zones can be defined by (domain, port) and/or by WWN zoning. To perform zoning to the granularity of the virtual N_Port IDs, you must use WWN-based zoning.

If you are using (domain, port) zoning for an NPIV port, and all the virtual PIDs associated with the port are included in the zone, then a port login (PLOGI) to a non-existent virtual PID is not blocked by the switch; rather, it is delivered to the device attached to the NPIV port. In cases where the device is not capable of handling such unexpected PLOGIs, you should use WWN-based zoning.

Enabling and Disabling NPIV

For Bloom-based switches and port blades supporting FICON (SilkWorm 3900 and 24000), NPIV is disabled for every port. For Condor-based switches and port blades (SilkWorm 4100, 4900, and 48000), NPIV is enabled for every port.

To enable or disable NPIV on a port-by-port basis, enter the **portCfgNPIVPort** command.

The following example shows NPIV being enabled on port 10 of a SilkWorm 4100:

```
switch:admin> portCfgNPIVPort 10, 1
```

Configuring NPIV

To specify the number of virtual N_Port_IDs per port or per switch, use the **configure** command with either of the following parameters:

- *switch.login.perPortMax*

  Use this parameter to set the number of virtual N_Port_IDs per port to a value between 0 and 255. The default setting is 126.

  For the SilkWorm 48000 with an FC4-48 port blade: For ports 0 through 15 on the FC4-48 port blade, the maximum number of virtual N_Port_IDs per port is 255; for ports 16 through 47, the maximum number is 127.

- *switch.login.perSwitchMax*

  Use this parameter to set the number of virtual N_Port_IDs per switch to a value between 0 and (126 * number of ports). The default setting is (15 * number of ports).

The following example shows the configuration of these parameters:

```
switch:admin> switchdisable
switch:admin> configure

Configure...

  Fabric parameters (yes, y, no, n): [no]
  Virtual Channel parameters (yes, y, no, n): [no]
  F-Port login parameters (yes, y, no, n): [no] y

    Maximum logins per switch: (1..4032) [4032] 2048
   Maximum logins per port: (1..255) [255] 126
.
.
.
switch:admin> switchenable
```

Configuration Scenarios

The actual number of virtual N_Port_IDs accepted per port and per switch is determined by the limits you set, and also by the limit that is reached *first*.

For example, if you've set *switch.login.perPortMax* to 25, and *switch.login.perSwitchMax* to 100, then the first 4 ports will accept 25 virtual N_Port_IDs each. However, the fifth port will reject any more virtual N_Port_IDs because the *switch.login.perSwitchMax* parameter (100) has been reached.

If you set *switch.login.perPortMax* to 50, and *switch.login.perSwitchMax* to 25, then the first port requests 50 virtual N_Port_IDs; 25 are accepted, and 25 are rejected. The second port requests 50 virtual N_Port_IDs, but none are accepted, because the *switch.login.perSwitchMax* limit of 25 has been reached.

Viewing NPIV Port Configuration Information

To view the NPIV capability of switch ports, enter the **portCfgShow** command. The following example shows whether or not a port is configured for NPIV:

```
switch:admin> portcfgshow
Ports of Slot 0    0  1  2  3    4  5  6  7    8  9 10 11   12 13 14 15
-----------------+--+--+--+--+----+--+--+--+----+--+--+--+----+--+--+--
Speed             AN AN AN AN   AN AN AN AN   AN AN AN AN   AN AN AN AN
Trunk Port        ON ON ON ON   ON ON ON ON   ON ON ON ON   ON ON ON ON
Long Distance     .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
VC Link Init      .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Locked L_Port     .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Locked G_Port     .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Disabled E_Port   .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
ISL R_RDY Mode    .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
RSCN Suppressed   .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Persistent Disable.. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
NPIV capability   ON ON ON ON   ON ON ON ON   ON ON ON ON   ON ON ON ON
```

Use the **switchShow** and **portShow** commands to view NPIV information for a given port. If a port is an F_Port, and you enter the **switchShow** command, then the port WWN of the N_Port is returned. For an NPIV F_Port, there are multiple N_Ports, each with a different port WWN. The **switchShow** command output indicates whether or not a port is an NPIV F_Port, and identifies the number of virtual N_Ports behind it. Following is sample output from the **switchShow** command:

```
switch: admin> switchshow
switchName:swd77
switchType:32.0
switchState: Online
switchMode:Native
switchRole:Principal
switchDomain: 99
switchId:fffc63
switchWwn:10:00:00:05:1e:35:37:40
zoning:      OFF
switchBeacon:OFF

Area Port Media Speed State
==============================
  0   0   id    N2   Online    F-Port  50:05:07:64:01:20:73:b8
  1   1   id    N2   Online    F-Port  50:05:07:64:01:60:73:b8
  2   2   id    N2   Online    F-Port  65 NPIV public
  3   3   id    N2   Online    F-Port  50:05:07:64:01:e0:73:b8
  4   4   id    N2   Online    F-Port  50:05:07:64:01:20:73:b5
  ...
  <output truncated>
```

Use the **portShow** command to view the NPIV attributes and all the N_Port (physical and virtual) port WWNs under “portWwn of device(s) connected.” Following is sample output for **portShow**:

```
switch:admin> portshow 2
portName: 02
portHealth: HEALTHY

Authentication: None
portDisableReason: None
portCFlags: 0x1
portFlags: 0x24b03 PRESENT ACTIVE F_PORT G_PORT NPIV LOGICAL_ONLINE LOGIN NOELP LED
ACCEPT
portType:  10.0
portState: 1Online
portPhys:  6In_Sync
portScn:   32F_Port
port generation number:    148
portId:    630200
portIfId:    43020005
portWwn:   20:02:00:05:1e:35:37:40
portWwn of device(s) connected:
    c0:50:76:ff:fb:00:16:fc
    c0:50:76:ff:fb:00:16:f8
    ...
      <output truncated>
      ...
    c0:50:76:ff:fb:00:16:80
    50:05:07:64:01:a0:73:b8
Distance:  normal
portSpeed: N2Gbps

Interrupts:        0          Link_failure: 16         Frjt:          0
Unknown:           0          Loss_of_sync: 422        Fbsy:          0
Lli:               294803     Loss_of_sig:  808
Proc_rqrd:         0          Protocol_err: 0
Timed_out:         0          Invalid_word: 0
Rx_flushed:        0          Invalid_crc:  0
Tx_unavail:        0          Delim_err:    0
Free_buffer:       0          Address_err:  1458
Overrun:           0          Lr_in:        15
Suspended:         0          Lr_out:       17
Parity_err:        0          Ols_in:       16
2_parity_err:      0          Ols_out:      15
CMI_bus_err:       0
```

Displaying Login Information

Use the **portLoginShow** command to display the login information for the virtual PIDs of a port. Following is sample output from **portLoginShow**:

```
switch:admin> portloginshow 2
Type  PID     World Wide Name        credit df_sz cos
=====================================================
  fe  630240  c0:50:76:ff:fb:00:16:fc   101  2048   c  scr=3
  fe  63023f  c0:50:76:ff:fb:00:16:f8   101  2048   c  scr=3
  fe  63023e  c0:50:76:ff:fb:00:17:ec   101  2048   c  scr=3
  ...
  <output truncated>
  ...
  ff  630202  c0:50:76:ff:fb:00:17:70   192  2048   c  d_id=FFFFFC
  ff  630201  c0:50:76:ff:fb:00:16:80   192  2048   c  d_id=FFFFFC
```

Section

II

Licensed Features

This section describes optionally licensed Brocade Fabric OS features and includes the following chapters:

- Chapter 17, “Administering Advanced Performance Monitoring”
- Chapter 18, “Administering Extended Fabrics”
- Chapter 19, “Administering ISL Trunking”
- Chapter 20, “Administering Advanced Zoning”
- Chapter 21, “Configuring and Monitoring FCIP Tunneling”
- Chapter 22, “Administering the iSCSI Target Gateway”

Additional Brocade licensed features are documented in the following books:

- *Fabric Watch Administrator’s Guide*
- *Secure Fabric OS Administrator’s Guide*
- *Fabric Manager Administrator’s Guide*

Administering Advanced Performance Monitoring

Chapter 17

This chapter contains information about the Brocade Advanced Performance Monitoring licensed feature, and contains the following topics:

- "About Advanced Performance Monitoring" next
- "Displaying and Clearing the CRC Error Count" on page 17-3
- "Monitoring End-to-End Performance" on page 17-3
- "Monitoring ISL Performance" on page 17-12
- "Monitoring Trunks" on page 17-12
- "Displaying Monitor Counters" on page 17-13
- "Clearing Monitor Counters" on page 17-15
- "Saving and Restoring Monitor Configurations" on page 17-16
- "Collecting Performance Data" on page 17-17

About Advanced Performance Monitoring

Based on Brocade Frame Filtering technology and a unique performance counter engine, Advanced Performance Monitoring is a comprehensive tool for monitoring the performance of networked storage resources. It supports direct-attach, loop, and switched fabric Fibre Channel SAN topologies by:

- Monitoring transaction performance from source to destination.
- Reporting cyclic redundancy check (CRC) error measurement statistics.
- Measuring Brocade InterSwitch Link (ISL) Trunking performance and resource usage.

Note

Advanced Performance Monitoring is not supported on VE_Ports (virtual FC_Ports) and EX_Ports. If you issue commands for any Advanced Performance Monitors on VE_Ports or EX_Ports you will receive error messages. See Chapter 21, "Configuring and Monitoring FCIP Tunneling" for more information about VE_Ports.

Additional features are provided through Web Tools. For additional information, see *Web Tools Administrator's Guide*.

Note

The SilkWorm 200E, 4100, 4900, 7500, and 48000 switches do not display AL_PA measurements for end-to-end monitors. CRC reports are accessed using Web Tools.

Also, AL_PA monitoring is not supported on NPIV configured ports.

Table 17-1 lists commands associated with Advanced Performance Monitoring. Advanced Performance Monitor commands are available only to users with the admin or switchAdmin roles. For detailed information on these commands, see the *Fabric OS Command Reference*.

Table 17-1 Advanced Performance Monitoring Commands

Command	Description
perfAddEEMonitor	Add an end-to-end monitor to a port.
perfAddIPMonitor	Add an IP monitor to a port.
perfAddReadMonitor	Add a SCSI Read monitor to a port.
perfAddRwMonitor	Add a SCSI Read and Write monitor to a port.
perfAddSCSIMonitor	Add a SCSI traffic frame monitor to a port.
perfAddUserMonitor	Add a filter-based monitor to a port.
perfAddWriteMonitor	Add a SCSI Write monitor to a port.
perfCfgClear	Clear the performance monitoring settings from nonvolatile (flash) memory.
perfCfgRestore	Restore performance monitoring settings from nonvolatile (flash) memory.
perfCfgSave	Save the current performance monitoring settings to nonvolatile (flash) memory.
perfClearAlpaCrc	Clear an AL_PA device CRC count by the port and AL_PA.
perfClearEEMonitor	Clear end-to-end monitor counters on a port.
perfClearFilterMonitor	Clears filter-based monitor counters.
perfDelEEMonitor	Delete an end-to-end monitor on port.
perfDelFilterMonitor	Delete a filter-based monitor.
perfMonitorClear	Clear statistics counters of end-to-end, filter-based, and ISL monitors on a port.
perfMonitorShow	Display end-to-end, filter-based, and ISL monitors on a port.
perfSetPortEEMask	Set overall mask for end-to-end (EE) monitors.
perfShowAlpaCrc	Display the AL_PA CRC count by port or by AL_PA.
perfShowEEMonitor	Show user-defined end-to-end monitors.
perfShowFilterMonitor	Show filter-based monitors.
perfShowPortEEMask	Display the current end-to-end mask of a port.

Note

The command examples in this chapter use the slot/port syntax required by SilkWorm 24000 and 48000 directors. For SilkWorm 200E, 3250, 3850, 3900, 4100, 4900 and 7500 switches, use only the port number where needed in the commands.

Displaying and Clearing the CRC Error Count

You can use the **perfShowAlpaCrc** command to display the CRC error count for all AL_PA devices or for a single AL_PA on a specific active L_Port.

The following example displays the CRC error count for all AL_PA devices on a port:

```
switch:admin> perfshowalpacrc 1/1
AL_PA    CRC count
--------------------
0xd9               0
```

The following example displays the CRC error count for a single AL_PA device on a port:

```
switch:admin> perfshowalpacrc 1/1, 0xd9
The CRC count at ALPA 0xd9 on port 1 is 0x000000000.
```

The following example clears the CRC error count:

```
switch:admin> perfclearalpacrc 1/1, 0xd9
CRC error count at AL_PA 0xd9 on port 1 is cleared.
switch:admin> perfclearalpacrc 1/1
No AL_PA value is specified. This will clear all AL_PA CRC
counts on port 1. Do you want to continue? (yes, y, no, n): [no] y
Please wait ...
All alpa CRC counts are cleared on port 1.
```

In v3.1.0, v4.1.0, and later versions, you can use the **portStatsClear** command clears AL_PA- based CRC error counters for all the ports in the same group.

Monitoring End-to-End Performance

End-to-end performance monitoring counts the number of words and CRC errors in Fibre Channel frames for a specified Source ID (SID) and Destination ID (DID) pair. An end-to-end performance monitor includes these counts:

- RX_COUNT (words in frames received at the port)
- TX_COUNT (words in frames transmitted from the port)
- CRC_COUNT (frames with CRC errors received at or transmitted from the port)

To enable end-to-end performance monitoring, you must configure an end-to-end monitor on a port, specifying the SID-DID pair (in hexadecimal). The monitor counts only those frames with matching SID and DID.

Each SID or DID has three fields, listed in the following order:

- Domain ID (DD)
- Area ID (AA)
- AL_PA (PP)

For example, the SID 0x118a0f denotes DD 0x11, AA 0x8a, and AL_PA 0x0f.

You can monitor end-to-end performance using the **perfMonitorShow** command, as described in "Displaying Monitor Counters" on page 17-13. You can clear end-to-end counters using the **perfMonitorClear** command, as described in "Clearing Monitor Counters" on page 17-15.

Note

For end-to-end monitors, CRC counters are not displayed on the SilkWorm 200E, 4100, 4900, 7500 and SilkWorm 48000 switches.

Adding End-to-End Monitors

An end-to-end monitor counts the following items for a port: number of words received, number of words transmitted, and number of CRC errors detected in frames.

SilkWorm 200E, 3250, 3850, 3900, and 24000 models allow up to eight end-to-end monitors.

The SilkWorm 4100, 4900, 7500 and 48000 models allow up to 256 end-to-end monitors shared by all ports in the same ASIC chip. (The number of interswitch links configured on the switch affects the amount of resources available for end-to-end monitors.)

The SilkWorm FR4-48 blade allows end-to-end monitors on all 48 ports.

For the FR4-16IP, end-to-end monitors are supported on the ASIC FC ports (ports 0 through 7), but not on the GbE ports.

You cannot add end-to-end monitors to interswitch links. Identical end-to-end monitors cannot be added to the same port. Two end-to-end monitors are considered identical if they have the same SID and DID values after applying the end-to-end mask.

The monitor count is qualified using either of following conditions:

- For frames received at the port with the end-to-end monitor installed, the frame SID is the same as "SourceID" and the frame DID is the same as "DestID". The RX_COUNT and CRC_COUNT are updated accordingly.
- For frames transmitted from the port with the end-to-end monitor installed, the frame DID is the same as "SourceID" and the frame SID is the same as "DestID". The TX_COUNT and CRC_COUNT are updated accordingly.

Note

The relationship between the area ID for a port and the port number depends on the PID format used by the fabric. See "Configuring the PID Format" on page A-1 for more information.

Figure 17-1 shows two devices:

- Host A is connected to domain 5 (0x05), switch area ID 18 (0x12), AL_PA 0x00 on Switch X
- Dev B is a storage device connected to domain 17 (0x11), switch area ID 30 (0x1e), AL_PA 0xef on Switch Y.

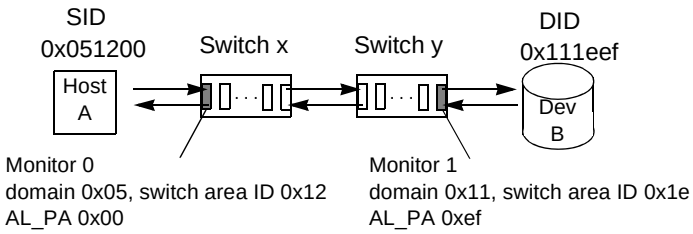

Figure 17-1 Setting End-to-End Monitors on a Port

Note

End-to-end performance monitoring looks at traffic on the receiving port respective to the SID only. In Figure 17-1, if you add a monitor to slot 2, port 2 on Switch x, specifying Dev B as the SID and Host A as the DID, no counters (except CRC) will be incremented.

To monitor the traffic from Host A to Dev B

```
switch:admin> perfaddeemonitor 2/2, "0x051200" "0x111eef"
End-to-End monitor number 0 added.
```

Add Monitor 0 to slot 2, port 2 on Switch x, specifying 0x051200 as the SID and 0x111eef as the DID, as shown in the following example:

Monitor 0 counts the frames that have an SID of 0x051200 and a DID of 0x111eef. For monitor 0, RX_COUNT is the number of words from Host A to Dev B, TX_COUNT is the number of words from Dev B to Host A, and CRC_COUNT is the number of frames in both directions with CRC errors.

To monitor the traffic from Dev B to Host A

```
switch:admin> perfaddeemonitor 2/14, "0x111eef" "0x051200"
End-to-End monitor number 1 added.
```

Add Monitor 1 to slot 2, port 14 on Switch y, specifying 0x111eef as the SID and 0x051200 as the DID, as shown in the following example.

Monitor 1 counts the frames that have an SID of 0x111eef and a DID of 0x051200. For monitor 1, RX_COUNT is the number of words from Dev B to Host A, TX_COUNT is the number of words from Host A to Dev B, and CRC_COUNT is the number of frames in both directions with CRC errors.

Figure 17-2 shows several switches and the correct ports on which to add performance monitors for a specified SID-DID pair.

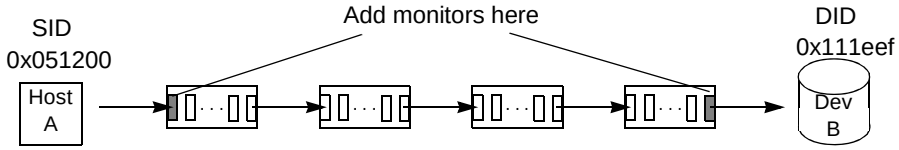

Figure 17-2 Proper Placement of End-to-End Performance Monitors

Setting a Mask for End-to-End Monitors

End-to-end monitors count the number of words in Fibre Channel frames that match a specific SID/DID pair. If you want to match only part of the SID or DID, you can set a mask on the port to compare only certain parts of the SID or DID. By default, the frame must match the entire SID and DID to trigger the monitor. By setting a mask, you can choose to have the frame match only one or two of the three fields (Domain ID, Area ID, and AL_PA) to trigger the monitor.

Note

Only one mask per port can be set. When you set a mask, all existing end-to-end monitors are deleted.

You can specify a mask using the **perfSetPortEeMask** command in the form dd:aa:pp, where dd is the domain ID mask, aa is the area ID mask, and pp is the AL_PA mask. The values for dd, aa, and pp are either ff (the field must match) or 00 (the field is ignored). The default EE mask value is ff:ff:ff. The **perfSetPortEeMask** command sets the mask for all end-to-end monitors of a port. If any end-to-end monitors are programmed on a port when the **perfSetPortEeMask** command is issued, then a message displays similar to the following example:

```
switch:admin> perfsetporteemask 1/2, "00:00:ff"
EE monitors are currently programmed on this port. Changing EE mask
for this port will cause ALL EE monitors on this port to be deleted.
Do you want to continue? (yes, y, no, n): [no] y

EE mask on port <port-number> is set and EE monitors were deleted
```

The **perfSetPorteeMask** command sets a mask for the Domain ID, Area ID, and AL_PA of the SIDs and DIDs for frames transmitted from and received by the port.

Figure 17-3 shows the mask positions in the command. A mask ("ff") is set on slot 1, port 2 to compare the AL_PA fields on the SID and DID in all frames (transmitted and received) on port 2. The frame SID and DID must match only the AL_PA portion of the specified SID-DID pair. Each port can have only one EE mask. The mask is applied to all end-to-end monitors on the port. Individual masks for each monitor on the port cannot be specified.

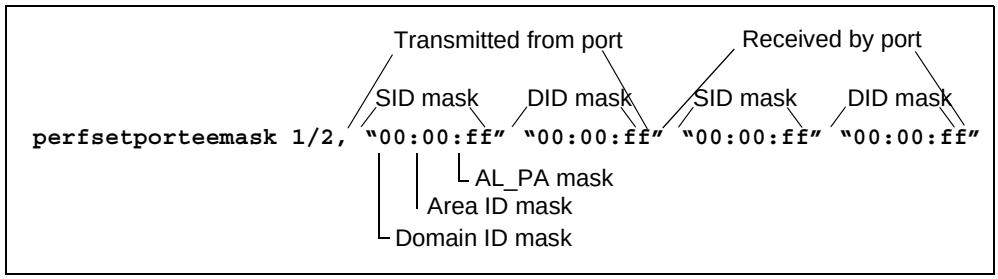

Figure 17-3 Mask Positions for End-to-End Monitors

To display the current end-to-end mask of a port

1. Enter the **perfShowPortEeMask** command.

The end-to-end mask has 12 fields, each with a value of on or off. See the following example of setting and displaying an end-to-end mask:

```
switch:admin> perfsetporteemask 1/11, "00:00:ff" "00:00:ff" "00:00:ff" "00:00:ff"
The EE mask on port 11 is set and EE counters are reset.
switch:admin> perfshowporteemask 1/11
The EE mask on port 11 is set by application TELNET
TxSID Domain:  off
TxSID Area:    off
TxSID AL_PA:    on
TxDID Domain:  off
TxDID Area:    off
TxDID AL_PA:    on
RxSID Domain:  off
RxSID Area:    off
RxSID AL_PA:    on
RxDID Domain:  off
RxDID Area:    off
RxDID AL_PA:    on
```

The end-to-end mask is set on slot 1, port 11.

The default EE mask value for an E_Port is:

- TxSID = "00:00:00"
- TxDID = "ff:00:00"
- RxSID = "00:00:00"
- RxDID = "00:00:00"

Only the transmit domain is monitored for E_Ports.

To display a monitor

1. Enter the **perfMonitorShow** command, as described in "Displaying Monitor Counters" on page 17-13.

Deleting End-to-End Monitors

Enter the **perfDelEeMonitor** command to delete end-to-end monitors. You can delete all monitors or specific monitors. The following example deletes the end-to-end monitor number 0 on slot 1, port 2:

```
switch:admin> perfdeleemonitor 1/2, 0
End-to-End monitor number 0 deleted
```

Monitoring Filter-Based Performance

Filter-based performance monitoring counts the number of times a frame with a particular pattern is transmitted by a port. Filter-based monitoring is achieved by configuring a filter for a particular purpose. The filter can be a standard filter (for example, a SCSI read command filter that counts the number of SCSI read commands that have been transmitted by the port) or a user-defined filter customized for your particular use.

For SilkWorm 200E, 3250, 3850, 3900, and 24000 models, the maximum number of filters is 8 per port, in any combination of standard filters and user-defined filters.

For the SilkWorm 4100, 4900, and 7500 models, the maximum number of filters is 12 per port, in any combination of standard filters and user-defined filters.

For the SilkWorm 48000 model, the maximum number of filters is 12 per port, in any combination of standard filters and user-defined filters, except for the FC4-48 port blade. For the FC4-48 port blade:

- Ports 0 through 15 have a maximum of 12 filter monitors per port.
- Ports 16 through 31 have a maximum of 6 filter monitors per port.
- Ports 32 through 47 do not have filter monitors.

For the SilkWorm FC4-16IP port blade, the maximum number of filters is 12 per port and 15 offsets per port.

The actual number of filters that can be configured on a port depends on the complexity of the filters. For trunked ports, the filter is configured on the trunk master.

You can monitor filter-based performance using the **perfMonitorShow** command, as described in "Displaying Monitor Counters" on page 17-13. You can clear filter-based counters using the **perfMonitorClear** command, as described in "Clearing Monitor Counters" on page 17-15.

Adding Standard Filter-based Monitors

Table 17-2 lists the commands for adding standard filter-based monitors to a port.

Table 17-2 Commands to Add Filter-Based Monitors

Telnet command	Description
perfAddReadMonitor	Count the number of SCSI Read commands
perfAddWriteMonitor	Count the number of SCSI Write commands
perfAddRwMonitor	Count the number of SCSI Read and Write commands
perfAddScsiMonitor	Count the number of SCSI traffic frames
perfAddIpMonitor	Count the number of IP traffic frames

You cannot add identical filter monitors to the same port. Two filter monitors are considered to be identical when they have the same:

- Filter monitor type
- Owner (telnet, Web Tools, etc.)
- Alias

The following example adds filter-based monitors to slot 1, port 2 and displays the results:

```
switch:admin> perfaddreadmonitor 1/2
SCSI Read filter monitor #0 added
switch:admin> perfaddwritemonitor 1/2
SCSI Write filter monitor #1 added
switch:admin> perfaddrwmonitor 1/2
SCSI Read/Write filter monitor #2 added
switch:admin> perfaddscsimonitor 1/2
SCSI traffic frame monitor #3 added
switch:admin> perfaddipmonitor 1/2
IP traffic frame monitor #4 added
switch:admin> perfmonitorshow --class FLT 1/2
There are 5 filter-based monitors defined on port 2.

KEY   ALIAS   OWNER_APP       OWNER_IP_ADDR     FRAME_COUNT
-----------------------------------------------------------------
0  SCSI Read  TELNET               N/A       0x0000000000000000
1  SCSI Write TELNET               N/A       0x0000000000000000
2  SCSI R/W   TELNET               N/A       0x0000000000000000
3  SCSI Frame TELNET               N/A       0x0000000000000000
4  IP Frame   TELNET                  N/A       0x0000000000000000
```

Adding Custom Filter-Based Monitors

In addition to the standard filters—read, write, read/write, SCSI frame and IP frame—you can create custom filters to gather statistics that fit your needs.

To define a custom filter, use the **perfAddUserMonitor** command. With this command, you must specify a series of *offsets*, *masks*, and *values*. For all transmitted frames, the switch performs these tasks:

- Locates the byte found in the frame at the specified *offset.*
- Applies the *mask* to the byte found in the frame.
- Compares the value with the given *values* in the **perfAddUserMonitor** command.
- Increments the filter counter if a match is found.

The following number of offsets can be specified:

Note

Bloom-based switches provide 2 unique unused offsets per port for user-defined filter monitors other than the 6 offsets already reserved for pre-defined filter monitors. Condor-based switches provide 15 unused offsets per port.

To illustrate, if you were to enter:

```
perfaddusermonitor 30 "4, 0xff, 0x06; 6, 0xff, 0x02; 7, 0xff, 0xd4"
```

the operation would be successful because offset 4 is a canned monitor and has a resource already reserved for it. Offsets 6 and 7 would exhaust all unique filter monitor resources on port 30. Therefore, any additional filter monitors created on port 30 would have to be canned filter monitors.

- SilkWorm 3016, 3250, 3850, 3900 and 24000 models (Fabric OS v4.0.0 or later)

  Up to two different offsets per port (one offset when FICON management server mode (FMS) is enabled).

- SilkWorm 4900, 7500 and 48000 model (Fabric OS v5.1.0 or later)

  Up to 15 different offsets per port (14 offsets when fmsmode is enabled).

- SilkWorm 3200, 3800 models (Fabric OS v3.0.0 or later)

  Up to three different offsets per port.

- SilkWorm 4100 model (Fabric OS v4.4.0 or later)

  Up to 15 different offsets per port (14 offsets when FMS is enabled).

- SilkWorm 200E and 4012 models (Fabric OS v5.0.1)

  Up to 7 different offsets per port (6 offsets when FMS is enabled).

You can specify up to four values to compare against each offset. If more than one offset is required to properly define a filter, the bytes found at each offset must match one of the given values for the filter to increment its counter. If one or more of the given offsets does not match any of the given values, the counter does not increment. The value of the offset must be between 0 and 63, in decimal format. Byte 0 indicates the first byte of the Start of Frame (SOF), byte 4 is the first byte of the frame header, and

byte 28 is the first byte of the payload. Thus only the SOF, frame header, and first 36 bytes of payload can be selected as part of a filter definition. Offset 0 is a special case, which can be used to monitor the first 4 bytes of the frame (SOF). When the offset is set to 0, the values 0–7 that are checked against that offset are predefined as shown in Table 17-3.

Table 17-3 Predefined Values at Offset 0

Value	SOF	Value	SOF
0	SOFf	4	SOFi2
1	SOFc1	5	SOFn2
2	SOFi1	6	SOFi3
3	SOFn1	7	SOFn3

If the switch does not have enough resources to create a given filter, then other filters might have to be deleted to free resources.

The following example adds filter-based monitors:

```
switch:admin> perfaddusermonitor 4/2, "12, 0xff, 0x05, 0x08; 9, 0xff, 0x02" "FCP/IP"
User monitor #5 added
switch:admin> perfaddusermonitor 1/2, "0, 0xff, 6"
User Monitor #6 added
```

Two filter-based monitors are added. The first monitor (#5) counts all FCP and IP frames transmitted from domain 0x02 for slot 4, port 2. The FCP and IP protocols are selected by monitoring offset 12, mask 0xff and matching values of 0x05 or 0x08. Domain 2 is selected by monitoring offset 9, mask 0xff, and matching a value of 0x02. The monitor counter is incremented for all outgoing frames from port 2 where byte 9 is 0x02 and byte 12 is 0x05 or 0x08.

The second monitor (#6) is for SOFi3 on slot 1, port 2.

Deleting Filter-Based Monitors

To delete a filter-based monitor

1. Entering the **perfShowFilterMonitor** command to list the valid monitor numbers.
2. Enter the **perfDelFilterMonitor** command to delete a specific monitor.

   If you do not specify which monitor number to delete, you are asked if you want to delete all entries.

The following example displays the monitors on slot 1, port 4 using the **perfShowFilterMonitor** command (the monitor numbers are listed in the KEY column) and deletes monitor number 1 on slot 1, port 4 using the **perfDelFilterMonitor** command:

```
switch:admin> perfshowfiltermonitor 1/4
There are 4 filter-based monitors defined on port 4.
KEY   ALIAS   OWNER_APP       OWNER_IP_ADDR     FRAME_COUNT
-----------------------------------------------------------------
0  SCSI Read  TELNET                N/A       0x0000000000002208
1  SCSI Write TELNET                N/A       0x000000000000464a
2  SCSI R/W   TELNET                N/A       0x000000000000fd8c
3  SCSI Frame WEB_TOOLS        192.168.169.40 0x00000000002c2229
switch:admin> perfdelfiltermonitor 1/4, 1
The specified filter-based monitor is deleted.
```

Monitoring ISL Performance

ISL monitoring is set up on E_Ports automatically in release v4.4.0 and later.

An ISL monitor measures traffic to all reachable destination domains for an ISL, showing which destination domain is consuming the most traffic. If there are more than 16 domains, the monitor samples traffic and extrapolates the measurement.

EE monitors on E_Ports are deleted when they become part of an ISL.

You can monitor ISL performance using the **perfMonitorShow** command, as described in "Displaying Monitor Counters." You can clear ISL counters using the **perfMonitorClear** command, as described in "Clearing Monitor Counters" on page 17-15.

Monitoring Trunks

For trunked ISLs on Fabric OS v4.x or higher switches, monitoring is set only on the master ISL, which communicates with the associated slave ISLs. Note the following:

- For Fabric OS v3.x switches, monitoring can be set on slave ISLs.
- End-to-end monitors are not supported for ISLs.
- SilkWorm 200E, 3250, 3850, 3900, and 24000 models support 8 filter-based monitors for trunks.
- The SilkWorm 4100, 4900, 7500, and 48000 switches support 12 filter-based monitors for trunks.

Displaying Monitor Counters

Use the **perfMonitorShow** command to display the monitors on a specified port. For end-to-end counters, you can display either the cumulative count of the traffic detected by the monitors or a snapshot of the traffic at specified intervals.

Note

SilkWorm 200E, 4100, 4900, 7500, and 48000 output does not include CRC counts.

The command format is:

```
perfmonitorshow --class monitor_class [slotnumber/]portnumber [interval]
```

monitor_class	The monitor class, which can be one of **EE** (end-to-end), **FLT** (filter-based), or **ISL** (inter-switch link). The **--class** *monitor_class* operand is required.
slotnumber	Specifies the slot number for a SilkWorm 24000 or 48000 director. For all other switches, this operand is not required. The slot number must be followed by a slash (/) and the port number, so that each port is represented by both slot number (1 through 4 or 7 through 10) and port number (0 through 15). The SilkWorm director has a total of 10 slots. Slot numbers 5 and 6 are control processor blades; slots 1 through 4 and 7 through 10 are port blades. For 16-port blades, there are 16 ports, counted from the bottom, numbered 0 to 15. For 32-port blades, there are 32 ports numbered 0 to 31.
portnumber	Specifies a port number. Valid values for port number vary, depending on the switch type. This operand is required.
interval	Specifies an interval in seconds. The interval must be greater than or equal to five seconds. For end-to-end monitoring, the Tx and Rx counts are measured in bytes. This operand is optional.

The following example displays end-to-end monitor on a port at 6-second intervals:

```
 switch:admin> perfMonitorShow --class EE 4/5 6
 perfmonitorshow 53, 6: Tx/Rx are # of bytes and crc is # of crc errors
     0            1            2            3            4
------------ ------------ ------------ ------------ ------------
crc    Tx    Rx crc    Tx    Rx crc    Tx    Rx crc    Tx    Rx crc    Tx    Rx
============ ============ ============ ============ ============
0     0      0    0     0     0    0     0      0    0     0     0    0     0    0
0    53m   4.9m  0    53m 4.9m 0    53m   4.9m  0    53m 4.9m  0    53m   0
0    53m   4.4m  0    53m 4.4m 0    53m   4.4m  0    53m 4.4m  0    53m   0
0    53m   4.8m  0    53m 4.8m 0    53m   4.8m  0    53m 4.8m  0    53m   0
0    53m   4.6m  0    53m 4.6m 0    53m   4.6m  0    53m 4.6m  0    53m   0
0    53m   5.0m  0    53m 5.0m 0    53m   5.0m  0    53m 5.0m  0    53m   0
0    53m   4.8m  0    53m 4.8m 0    53m   4.8m  0    53m 4.8m  0    53m   0
0    53m   4.5m  0    53m 4.5m 0    53m   4.5m  0    53m 4.5m  0    53m   0
0    52m   4.5m  0    52m 4.5m 0    52m   4.5m  0    52m 4.5m  0    52m   0
0    52m   5.0m  0    52m 5.0m 0    52m   5.0m  0    52m 5.0m  0    52m   0
0    52m   4.5m  0    52m 4.5m 0    52m   4.5m  0    52m 4.5m  0    52m   0
0    52m   4.6m  0    52m 4.6m 0    52m   4.6m  0    52m 4.6m  0    52m   0
```

The following example displays EE monitors on a port:

```
switch:admin> perfMonitorShow --class EE 4/5
There are 7 end-to-end monitor(s) defined on port 53.
KEY     SID     DID     OWNER_APP    OWNER_IP_ADDR      TX_COUNT          RX_COUNT           CRC_COUNT
-------------------------------------------------------------------------------------------------------
  0  0x58e0f  0x1182ef   TELNET          N/A        0x0000000000000000  0x0000000000000000  0x0000000000000000
  0   0x21300  0x21dda    TELNET          N/A         0x00000004d0ba9915  0x0000000067229e65 0x0000000000000000
  1  0x21300  0x21ddc    TELNET          N/A         0x00000004d0baa754  0x0000000067229e65 0x0000000000000000
  2  0x21300  0x21de0    TELNET          N/A         0x00000004d0bab3a5  0x0000000067229e87 0x0000000000000000
  3  0x21300  0x21de1    TELNET          N/A         0x00000004d0bac1e4  0x0000000067229e87 0x0000000000000000
  4  0x21300  0x21de2    TELNET          N/A         0x00000004d0bad086  0x0000000067229e87 0x0000000000000000
  5  0x11000  0x21fd6   WEB_TOOLS    192.168.169.40  0x00000004d0bade54  0x0000000067229e87 0x0000000000000000
  6  0x11000  0x21fe0   WEB_TOOLS    192.168.169.40  0x00000004d0baed41  0x0000000067229e98 0x0000000000000000
```

The following example displays filter-based monitor on a port at 6-second intervals:

```
switch:admin> perfMonitorShow --class FLT 2/5 6
perfmonitorshow 21, 6
    0        1        2        3        4        5        6
 #Frames  #Frames  #Frames  #Frames  #Frames  #Frames  #Frames
-----------------------------------------------------------------
   0        0        0        0        0        0        0
  26k     187      681      682      682      494      187
  26k     177      711      710      710      534      176
  26k     184      734      734      734      550      184
  26k     182      649      649      649      467      182
  26k     188      754      755      755      567      184
  26k     183      716      716      717      534      183
  26k     167      657      656      655      488      167
  26k     179      749      749      749      570      179
  26k     164      752      752      752      588      164
  26k     190      700      700      700      510      190
  26k     181      701      701      701      520      181
  26k     200      750      750      751      550      201
  26k     180      692      692      691      512      179
  26k     179      696      696      696      517      179
  26k     187      720      720      720      533      187
  26k     200      722      722      722      522      200
  26k     204      717      717      717      513      204
```

The following example displays filter monitor information on a port:

```
switch:admin> perfMonitorShow --class FLT 2/5
There are 7 filter-based monitors defined on port 21.

KEY   ALIAS   OWNER_APP       OWNER_IP_ADDR     FRAME_COUNT
-------------------------------------------------------------
 0  SCSI_Frame TELNET              N/A      0x00000000002c2229
 1  SCSI_WR    TELNET              N/A      0x000000000000464a
 2  SCSI_RW    TELNET              N/A      0x000000000000fd8c
 3  SCSI_RW   WEB_TOOLS      192.168.169.40  0x0000000000007ba3
 4  SCSI_RW   WEB_TOOLS      192.168.169.190 0x0000000000004f0e
 5  SCSI_RD   WEB_TOOLS      192.168.169.40  0x0000000000002208
 6  SCSI_WR   WEB_TOOLS      192.168.169.40  0x000000000000033a
```

The following example displays ISL monitor information on a port:

```
switch:admin> perfMonitorShow --class ISL 1/1
Total transmit count for this ISL: 1462326
Number of destination domains monitored: 3
Number of ports in this ISL: 2
Domain  97:              110379          Domain 98:              13965
Domain  99:             1337982
```

Clearing Monitor Counters

Before you clear statistics counters, verify the valid monitor numbers on a specific port using the **perfMonitorShow** command, to make sure the correct monitor counters are cleared. To clear statistics counters for all or a specified monitor, use the **perfMonitorClear** command. After the command has been executed, the telnet shell confirms that the counters on the monitor have been cleared.

The command format is:

```
perfmonitorclear --class monitor_class [slotnumber/]portnumber [monitorId]
```

monitor_class	The monitor class, which can be one of **EE** (end-to-end), **FLT** (filter-based), or **ISL** (inter-switch link). The **--class** *monitor_class* operand is required.
slotnumber	Specifies the slot number for a SilkWorm 24000 or 48000 director. For all other switches, this operand is not required. The slot number must be followed by a slash (/) and the port number, so that each port is represented by both slot number (1 through 4 or 7 through 10) and port number (0 through 15). The SilkWorm director has a total of 10 slots. Slot numbers 5 and 6 are control processor blades; slots 1 through 4 and 7 through 10 are port blades. For 16-port blades, there are 16 ports, counted from the bottom, numbered 0 to 15. For 32-port blades, there are 32 ports numbered 0 to 31.
portnumber	Specifies a port number. Valid values for port number vary, depending on the switch type. This operand is required.
monitorId	Specifies the monitor number to clear. Monitor numbers are defined when you create the monitor on a port. This operand is optional. If not specified, all monitor counters on the port are cleared. This operand does not apply to ISL monitors.

Note

In Fabric OS v3.1.0 and v4.1.0 (or later) the **portStatsClear** command clears AL_PA-based CRC error counters for all the ports in the same group.

The following example clears statistics counters for an end-to-end monitor:

```
switch:admin> perfMonitorClear --class EE 1/2 5
End-to-End monitor number 5 counters are cleared

switch:admin> perfMonitorClear --class EE 1/2
This will clear ALL EE monitors' counters on port 2, continue?
(yes, y, no, n): [no] y
```

The following example clears statistics counters for a filter-based monitor:

```
switch:admin> perfMonitorClear --class FLT 1/2 4
Filter-based monitor number 4 counters are cleared

switch:admin> perfMonitorClear --class FLT 1/2
This will clear ALL filter-based monitors' counters on port 2, continue? (yes, y, no,
y): [no] y
```

The following example clears statistics counters for an ISL monitor:

```
switch:admin> perfMonitorClear --class ISL 1
This will clear ISL monitor on port 1, continue? (yes, y, no, n): [no] y
```

Saving and Restoring Monitor Configurations

To save the current end-to-end and filter monitor configuration settings into nonvolatile memory, use the **perfCfgSave** command:

```
switch:admin> perfcfgsave
This will overwrite previously saved Performance Monitoring
settings in FLASH. Do you want to continue? (yes, y, no, n): [no] y
Please wait ...
Performance monitoring configuration saved in FLASH.
```

To restore a saved monitor configuration, use the **perfCfgRestore** command. For example, to restore the original performance monitor configuration after making several changes:

```
switch:admin> perfcfgrestore
This will overwrite current Performance Monitoring settings in RAM. Do you want to
continue? (yes, y, no, n): [no] y
Please wait... Performance monitoring configuration restored from FLASH ROM.
```

To clear the previously saved performance monitoring configuration settings from nonvolatile memory, use the **perfCfgClear** command:

```
switch:admin> perfcfgclear
This will clear Performance Monitoring settings in FLASH. The RAM settings won't
change.  Do you want to continue? (yes, y, no, n): [no] y
Please wait... Committing configuration...done.
Performance Monitoring configuration cleared from FLASH.
```

To prevent the switch configuration flash from running out of memory, the number of monitors saved to flash memory is limited as follows:

- The total number of EE monitors per port is limited to 16.
- The total number of filter monitors per port is limited to 16.
- The total number of monitors per switch is limited to 512.

When there are more than 512 monitors in the system, monitors are saved to flash memory in the following order:

- The EE monitors for each port (from 0 to MAX_PORT)
- The filter monitors for each port

So EE monitors get preference saving to flash memory when the total number of monitors in a switch exceeds 512. If the total number of monitors per port or switch exceeds the limit, then you will receive an error message indicating the count has been exceeded and that some monitors have been discarded.

Collecting Performance Data

Data collected through Advanced Performance Monitoring is deleted when the switch is rebooted. Using the Brocade Fabric Manager software application version 4.4.0 or later, you can store performance data persistently. For details on this feature, see the *Fabric Manager Administrator's Guide*.

Administering Extended Fabrics

Chapter 18

This chapter contains procedures for using the Brocade Extended Fabrics licensed feature, which extends the distance that interswitch links (ISLs) can reach over a dark fiber or wavelength-division multiplexing (WDM) connection. The Extended Fabrics feature is not used over FCIP connections over IP WANs. To use extended ISL modes, you must first install the Extended Fabrics license. For details on obtaining and installing licensed features, see "Maintaining Licensed Features" on page 2-14.

This chapter contains the following topics:

- "About Extended Link Buffer Allocation" on page 18-1
- "Fabric Considerations" on page 18-2
- "Choosing an Extended ISL Mode" on page 18-3
- "Configuring External Ports" on page 18-7
- "Configuring an Extended ISL" on page 18-10

About Extended Link Buffer Allocation

As the distance between switches and the link speed increases, additional *buffer-to-buffer credits* are required to maintain maximum performance. The number of credits reserved for a port depends on the switch model and on the extended ISL mode for which it is configured.

SilkWorm 3250, 3850, 3900, 24000, and 48000 (FC2-16 Port Blades)

For the following switches, each *port group* contains four ports and uses a common pool of credits.

- SilkWorm 3250
- SilkWorm 3850
- SilkWorm 3900
- SilkWorm 24000
- SilkWorm 48000 using FC2-16 port blades

Because the number of credits available for use within each port group is limited, configuring ports for extended links on these models might cause other ports to become disabled if there are not enough buffer credits available; for example:

- If two 2-Gbit/second ports in a group are configured for L1 mode, each will be allocated enough buffer-to-buffer credits to cause the other two ports in the group to become disabled.

- A port connected to a device that is in loopback mode might become disabled for lack of buffers if another port in that group is set to L2 mode.

See Chapter 9, “Configuring SilkWorm Directors” for details about port blade nomenclature.

SilkWorm 4100, 4900, 7500, and 48000 (FC4-16, FC4-16IP, FC4-32, FC4-48, and FR4-18i Port Blades)

For the following switches, buffer credits are used by all ports on the chip. *Buffer-limited* port technology allows all ports to remain operational, even when extended links are in use.

- SilkWorm 4100
- SilkWorm 4900
- SilkWorm 7500
- SilkWorm 48000 using the following port blades:
  - FC4-16
  - FC4-16IP
  - FC4-32
  - FC4-48
  - FR4-18i

A buffer-limited port can come online with a minimum of 8 buffer credits, which could be less than its configuration specifies, allowing the port to continue to operate rather than being disabled due to a lack of buffers. The 8 buffer minimum should allow 4 Gbit/sec speeds over distances within most data centers.

Buffer-limited operation is supported for the L0 and LD extended ISL modes only, and is persistent across reboots, switch disabling and enabling, and port disabling and enabling.

Fabric Considerations

Consider these items that affect the fabric when you configure extended ISLs:

- Balance the number of long-distance ISL connections and core-to-edge ISL connections within a switch. Configuring long-distance ISLs between core and edge switches is possible, but is not a recommended practice.
- Starting with Fabric OS v4.4.0, VC translation link initialization (an option of the **portCfgLongDistance** command) is enabled by default for long-distance links. For previous Fabric OS versions that support this option, it was disabled by default. To avoid inconsistency in the fabric, make sure that this value is enabled on both ends of the link. To connect to switches running Fabric OS versions earlier than v4.0.2 and v3.0.2c, make sure that VC translation link initialization is disabled because these versions do not support it.

Choosing an Extended ISL Mode

The following extended ISL modes are supported:

- L0—Level 0 static mode, the default.
- LE—Level E static mode, supports links beyond 5 km.
- LD—Dynamic mode uses automatic distance detection for a user-specified distance.
- LS—Static long distance mode uses the user-specified link distance.

The following extended ISL modes are supported; however, you cannot newly configure these modes on a Fabric OS 5.2.0 and higher switch. If you configure these modes on a switch running an earlier Fabric OS version and then upgrade the switch to Fabric OS 5.2.0 or higher, these modes continue to be supported.

- L0.5—Level 0.5 static mode (designated LM when listed with the **portcfgshow** command).
- L1—Level 1 static mode.
- L2—Level 2 static mode.

For dynamic long distance links, you can approximate the number of buffer credits using the following formula:

*Buffer credits = [(distance in km) * (data rate) * 1000] / 2112*

The data rate is 1.0625 for 1 Gbit/sec, 2.125 for 2 Gbit/sec, and 4.25 for 4 Gbit/sec and Fibre Channel.

Switches with Bloom ASIC

Table 18-1 lists the extended ISL modes for switches that have a Bloom ASIC. You can configure extended ISL modes with the **portCfgLongDistance** command when the Extended Fabrics license is activated.

Table 18-1 Extended ISL Modes: Switches with Bloom and Bloom II ASICs

Mode	Buffer Allocation		Distance @ 1 Gbit/sec	Distance @ 2 Gbit/sec	Earliest Fabric OS Release	Extended Fabrics License Required?
	1 Gbit/sec	2 Gbit/sec				
L0	5 (26)[a]	5 (26)	10 km	5 km	All	No
LE	13	19	n/a	10 km	v3.x, v4.x	No
L0.5[b]	19	34	25 km	25 km	v3.1.0, v4.1.0, v4.x, 5.x	Yes
L1[b]	27	54	50 km	50 km	All	Yes
L2[b]	60	65 108 for Bloom II	100 km	60 km 100 km for Bloom II	All	Yes
LD[c]	Auto	Auto	Auto Maximum is 200 km	Auto Maximum is 100 km	v3.1.0, v4.1.0, v4.4.0, 5.x (depending on the model)	Yes
LS[d]	Based on user-specified distance	Based on user-specified distance	Based on user-specified distance Maximum is 200 km	Based on user-specified distance Maximum is 100 km	v5.1.0	Yes

a. For each data channel (in this case, there are 4) there are 5 credits, plus 6 extra credits.

b. This mode is supported but cannot be configured on switches running Fabric OS 5.2.0 and higher.

c. The dynamic long-distance mode (LD) automatically configures the number of buffer credits required, based on the actual link distance.

d. The static long-distance mode (LS) allocates the number of buffer credits based on the user-specified distance.

Switches with Goldeneye ASIC

Table 18-2 lists the extended ISL modes for the SilkWorm 200E, which has a Goldeneye ASIC.

Table 18-2 SilkWorm 200E Extended ISL Modes (Goldeneye ASIC)

Mode[a]	Buffer Allocation			Distance @ 1 Gbit/sec	Distance @ 2 Gbit/sec	Distance @ 4 Gbit/sec	Earliest Fabric OS Release	Extended Fabrics License Required?
	1 Gbit/ sec	2 Gbit/ sec	4 Gbit/ sec					
L0	3(17)[a]	3(17)	3(17)	6 km	3 km	1.5 km	All	No
LE	11	16	26	10 km	10 km	10 km	v3.x, v4.x	No
L0.5[b]	18	31	56	25 km	25 km	25 km	v5.1.0	Yes
L1[b]	31	56	106	50 km	50 km	50 km	v5.1.0	Yes
L2[b]	56	106	n/a	100 km	100 km	n/a	v5.1.0	Yes
LD	Auto[c]	Auto	Auto	Auto	Auto	Auto	v5.1.0	Yes
LS[d]	Based on user-specified distance	Based on user-specified distance	Based on user-specified distance	Based on user-specified distance Maximum is 293 km.	Based on user-specified distance Maximum is 146 km.	Based on user-specified distance Maximum is 73 km.	v5.1.0	Yes

a. For each data channel (in this case there are 4) there are 3 credits, plus 5 extra credits.

b. This mode is supported but cannot be configured on switches running Fabric OS 5.2.0 and higher.

c. The dynamic long-distance mode (LD) automatically configures the number of buffer credits required, based on the actual link distance.

d. The static long-distance mode (LS) allocates the number of buffer credits based on the user-specified distance.

Switches with Condor ASIC

Table 18-3 lists the extended ISL modes for switches that have a Condor ASIC.

Table 18-3 Extended ISL Modes: Switches with Condor ASIC

Mode	Buffer Allocation			Distance @ 1 Gbit/sec	Distance @ 2 Gbit/sec	Distance @ 4 Gbit/sec	Earliest Fabric OS Release	Extended Fabrics License Required?
	1 Gbit/ sec	2 Gbit/ sec	4 Gbit/ sec					
L0	5 (26)[a]	5 (26)	5 (26)	10 km	5 km	2 km	All	No
LE	11	16	26	10 km	10 km	10 km	v3.x, v4.x	No
L0.5[b]	18	31	56	25 km	25 km	25 km	v3.1.0, v4.1.0, v4.x, 5.x	Yes
L1[b]	31	56	106	50 km	50 km	50 km	All	Yes
L2[b]	56	106	206	100 km	100 km	100 km	All	Yes
LD[c]	Auto	Auto	Auto	Auto Maximum is 500 km	Auto Maximum is 250 km	Auto Maximum is 100 km	v3.1.0, v4.1.0, v4.4.0, 5.x (depending on the model)	Yes
LS[d]	Based on user-specified distance	Based on user-specified distance	Based on user-specified distance	Based on user-specified distance Maximum is 500 km	Based on user-specified distance Maximum is 250 km	Based on user-specified distance Maximum is 100 km	v5.1.0	Yes

a. For each data channel (in this case, there are 4) there are 5 credits, plus 6 extra credits.

b. This mode is supported but cannot be configured on switches running Fabric OS 5.2.0 and higher.

c. The dynamic long-distance mode (LD) automatically configures the number of buffer credits required, based on the actual link distance.

d. The static long-distance mode (LS) allocates the number of buffer credits based on the user-specified distance.

Configuring External Ports

The number of ports that can be configured per port group for each switch depends on both port speed and distance. The following sections specify the number of external ports available for each switch at a specific port speed and distance.

SilkWorm 200E

The number of ports that can be configured per port group at various distances is summarized in Table 18-4.

Table 18-4 SilkWorm 200E

Speed (Gbit/sec)	Number of Ports Allowed at Distance (km)					
	10 km	25 km	50 km	100 km	250 km	500 km
1	16 ports	12 ports	6 ports	3 ports	1 port	n/a
2	16 ports	6 ports	3 ports	1 port	n/a	n/a
4	8 ports	3 ports	1 port	0	n/a	n/a

SilkWorm 3250, 3850, 3900, 24000, and 48000 (FC2-16 port blades)

The number of ports that can be configured per port group at various distances is summarized in Table 18-5.

Table 18-5 SilkWorm 3250, 3850, 3900, 24000, and 48000 (FC2-16 port blades)

Speed (Gbit/sec)	Number of Ports Allowed at Distance (km)					
	10 km	25 km	50 km	100 km	250 km	500 km
1	4 ports	4 ports	L1 mode: 4 ports LD mode: 3 ports	up to 2 ports	n/a	n/a
2	4 ports	3 ports	up to 2 ports	1port	n/a	n/a
4	n/a	n/a	n/a	n/a	n/a	n/a

SilkWorm 4100

The number of ports that can be configured at various distances is summarized in Table 18-6.

Table 18-6 SilkWorm 4100

Speed (Gbit/sec)	Number of Ports Allowed at Distance (km)					
	10 km	25 km	50 km	100 km	250 km	500 km
1	32 ports	32 ports	32 ports	up to 15 ports	up to 6 ports	up to 3 ports
2	32 ports	32 ports	up to 15 ports	up to 7 ports	up to 3 ports	n/a
4	32 ports	up to 15 ports	up to 7 ports	up to 3 ports	n/a	n/a

SilkWorm 4900

The number of ports that can be configured at various distances is summarized in Table 18-7.

Table 18-7 SilkWorm 4900

Speed (Gbit/sec)	Number of Ports Allowed at Distance (km)					
	10 km	25 km	50 km	100 km	250 km	500 km
1	64 ports	64 ports	64 ports	44 ports	16 ports	8 ports
2	64 ports	64 ports	48 ports	20 ports	8 ports	n/a
4	64 ports	48 ports	20 ports	8 ports	n/a	n/a

SilkWorm 7500

For the SilkWorm 7500, long distance settings are applicable only to the physical FC_Ports (ports 0-15). Long distance settings are not applicable to virtual FC_Ports (ports 16-31). The number of ports that can be configured at various distances is summarized in Table 18-8.

Table 18-8 SilkWorm 7500

Speed (Gbit/sec)	Number of Ports Allowed at Distance (km)					
	10 km	25 km	50 km	100 km	250 km	500 km
1	16 ports	16 ports	16 ports	16 ports	6 ports	2 ports
2	16 ports	16 ports	16 ports	8 ports	2 ports	n/a
4	16 ports	16 ports	8 ports	4 ports	n/a	n/a

SilkWorm 48000 (FC4-16 port blades)

The number of ports that can be configured at various distances is summarized in Table 18-9.

Table 18-9 SilkWorm 48000 (FC4-16 blades)

Speed (Gbit/sec)	Number of Ports Allowed at Distance (km)					
	10 km	25 km	50 km	100 km	250 km	500 km
1	16 ports	16 ports	16 ports	up to 12 ports	up to 4 ports	up to 2 ports
2	16 ports	16 ports	up to 12 ports	up to 5 ports	up to 2 ports	0
4	16 ports	up to 12 ports	up to 5 ports	up to 2 ports	0	0

SilkWorm 48000 (FC4-16IP blades)

The number of ports that can be configured at various distances is summarized in Table 18-10.

Table 18-10 SilkWorm 48000 (FC4-16IP blades)

Speed (Gbit/sec)	Number of Ports Allowed at Distance (km)					
	10 km	25 km	50 km	100 km	250 km	500 km
1	8 ports	8 ports	8 ports	8 ports	5 ports	2 ports
2	8 ports	8 ports	8 ports	6 ports	2 ports	0 port
4	8 ports	8 ports	6 ports	3 ports	0 port	0

SilkWorm 48000 (FC4-32 port blades)

The number of ports that can be configured at various distances is summarized in Table 18-11.

Table 18-11 SilkWorm 48000 (FC4-32 blades)

Speed (Gbit/sec)	Number of Ports Allowed at Distance (km)					
	10 km	25 km	50 km	100 km	250 km	500 km
1	32 ports	32 ports	32 ports	up to 26 ports	up to 10 ports	up to 4 ports
2	32 ports	32 ports	up to 26 ports	up to 12 ports	up to 4 ports	0
4	32 ports	up to 26 ports	up to 12 ports	up to 6 ports	0	0

SilkWorm 48000 (FC4-48 port blades)

The number of ports that can be configured at various distances is summarized in Table 18-12.

Table 18-12 SilkWorm 48000 (FC4-48 blades)

Speed (Gbit/sec)	Number of Ports Allowed at Distance (km)					
	10 km	25 km	50 km	100 km	250 km	500 km
1	48 ports	48 ports	48 ports	up to 22 ports	up to 8 ports	up to 4 ports
2	48 ports	48 ports	up to 22 ports	up to 10 ports	up to 4 ports	0
4	32 ports	up to 22 ports	up to 10 ports	up to 4 ports	0	0

SilkWorm 48000 (FR4-18i blades)

For the FR4-18i blade, long distance settings are applicable only to the physical FC_Ports (ports 0-15). Long distance settings are not applicable to virtual FC_Ports (ports 16-31). The number of ports that can be configured at various distances is summarized in Table 18-13.

Table 18-13 SilkWorm 48000 (FR4-18i blades)

Speed (Gbit/sec)	Number of Ports Allowed at Distance (km)					
	10 km	25 km	50 km	100 km	250 km	500 km
1	16 ports	16 ports	16 ports	16 ports	6 ports	2 ports
2	16 ports	16 ports	16 ports	8 ports	2 ports	n/a
4	16 ports	16 ports	8 ports	4 ports	n/a	n/a

Configuring an Extended ISL

Before configuring the ISL, ensure that the following conditions are met:

- Extended ISL support for SilkWorm 2000-series switches is limited as follows:
  - Extended ISLs are not supported between SilkWorm 2000-series switches and other SilkWorm models.
  - To support extended ISLs between Silkworm 2000-series switches, the **fabric.ops.mode.longDistance** parameter must be set to 1 on all switches in the fabric. Each switch must be disabled before setting this parameter.
- For fabrics that contain a mix of SilkWorm models, the **fabric.ops.mode.longDistance** parameter must be set to 0 (the default). Under certain circumstances (for example, if you want extended distance between Bloom-based switches) this mode needs to be enabled (set to 1) on switches running Fabric OS v3.x or v4.x. Talk to your switch provider for details.
- The ports on both ends of the ISL must have the same configuration.
- Use only qualified SFPs.

To configure an extended ISL

1. Connect to the switch and log in as admin.
2. If the fabric contains SilkWorm 2000-series extended ISLs, use the **switchDisable** command to disable the switch and then use the **configure** command to set the fabric-wide configuration parameter **fabric.ops.mode.longDistance** to 1 on all switches in the fabric.
3. Enter the **portCfgLongDistance** command, using the following syntax:

```
portcfglongdistance [slotnumber/]portnumber [distance_level]
[vc_translation_link_init] [desired_distance]
```

slotnumber	Specify the slot number for SilkWorm 24000 and 48000 directors. This option is not applicable to fixed-port switches. The slot number must be followed by a slash (/) and the port number.
portnumber	Specify the port number.
distance_level	Specify the ISL mode to be set on the port (see “Choosing an Extended ISL Mode” on page 18-3).

vc_translation_link_init	This extended link initialization sequence, which is an enhanced link reset protocol, avoids excessive resetting of ports. By default this option is set to 1 (enabled). To prevent fabric segmentation, this option must be set to 0 (disabled) when connecting to switches running Fabric OS versions earlier than v3.0.2c or v4.0.2. It must be set to 1 (enabled) when configuring a trunk over extended fabrics.
desired_distance	Required for a port configured for LD mode. Specify the desired distance, in kilometers, for the link. The specified value is the upper limit for calculating buffer availability for the port. If the measured distance is more than the specified *desired_distance*, the port is allocated the number of buffers required by the specified desired distance. (Fabric OS versions earlier than v4.4.0 do not support this parameter.) **Note:** For SilkWorm 48000 with FC4-16, FC4-16IP, FC4-32, or FC4-48 port blades, ports will be shown as buffer limited.

4. Repeat step 3 for the remote extended ISL port. Both the local and remote extended ISL ports must be configured to the same distance level. When the connection is initiated, the fabric will reconfigure.

The following example configures slot 1, port 1 for the LD link distance mode, enables the extended link initialization sequence, and sets the desired distance to 50 kilometers:

```
switch:admin> portcfglongdistance 1/1 LD 1 50
```

Administering ISL Trunking

Chapter 19

This chapter contains procedures for using the Brocade ISL Trunking licensed feature, which optimizes the use of bandwidth by allowing a group of interswitch links to merge into a single logical link.

This chapter contains the following topics:

- "About ISL Trunking" on page 19-1
- "Standard Trunking Criteria" on page 19-2
- "Fabric Considerations" on page 19-3
- "Initializing Trunking on Ports" on page 19-4
- "Monitoring Traffic" on page 19-5
- "Enabling and Disabling ISL Trunking" on page 19-6
- "Setting Port Speeds" on page 19-7
- "Displaying Trunking Information" on page 19-8
- "Trunking Over Extended Fabrics" on page 19-9
- "Troubleshooting Trunking Problems" on page 19-11

About ISL Trunking

Brocade ISL Trunking reduces or eliminates situations that require static traffic routes and individual ISL management to achieve optimal performance. Trunking optimizes fabric performance by distributing traffic across the shared bandwidth of all the interswitch links in a trunking group, allowing traffic to flow through any available link in a group rather than restricting it to a specific, potentially congested link. The use of trunking results in simplified fabric design and management, lowered cost of ownership, and increased data availability.

To use trunking, you must first install the Brocade ISL Trunking license. For details on obtaining and installing licensed features, see "Maintaining Licensed Features" on page 2-14. Trunking is enabled automatically when the ISL Trunking license is activated and ports are reinitialized (after installing the license, you enter the **switchDisable** and **switchEnable** commands), and trunks are easily managed using either Fabric OS CLI commands or Web Tools. You can enable and disable trunking and set trunk port speeds (for example, 2 Gig/sec, 4 Gig/sec, or autonegotiate) for entire switches or for individual ports.

Trunks distribute traffic dynamically and in order at the frame level, achieving greater performance with fewer interswitch links.

Trunks are compatible with both short wavelength (SWL) and long wavelength (LWL) fiber optic cables and transceivers.

Figure 19-1 illustrates how trunking can result in more throughput by distributing data over four ISLs with no congestion. In a fabric that does not have trunking capability, some paths would be congested and other paths under-utilized.

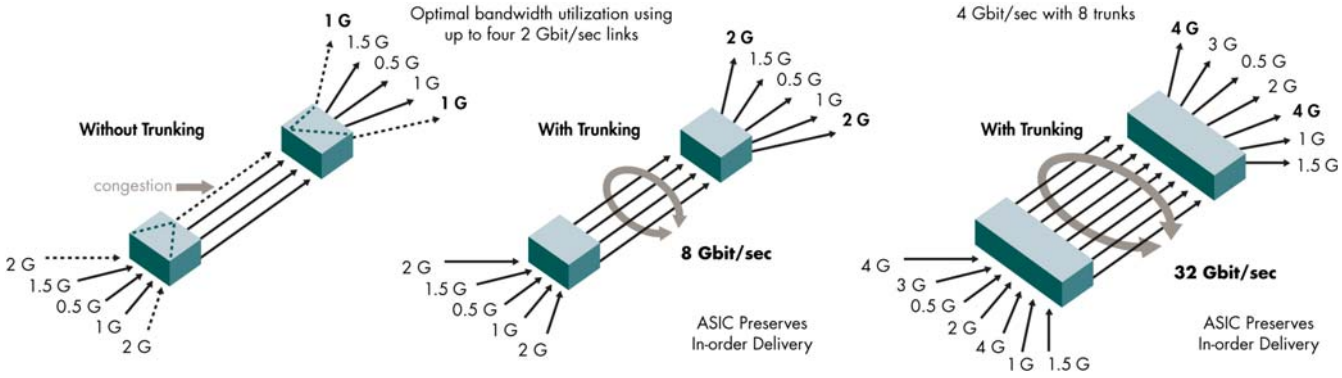

Figure 19-1 Distribution of Traffic over ISL Trunking Groups

Trunks operate best when the cable length of each trunked link is roughly equal to the others in the trunk. For optimal performance, no more than 30 meters difference is recommended.

Connections between SilkWorm 4100, 4900, and 48000 (using FC4-16, FC4-32, FC4-48, and FC4-16IP port blades) models support these advanced features:

- Up to eight ports in one trunk group to create high performance 32-Gbit/sec ISL trunks between switches
- ISL Trunking over longer distances than other models
- Dynamic trunk master reassignment if a trunk master is disabled (on other platforms, all ports on a trunk must be disabled temporarily to reassign a master)
- 4 Gbit/sec trunk links

The maximum number of ports per trunk and trunks per switch depends on the SilkWorm model.

For detailed information about trunking commands, see online help or *Fabric OS Command Reference*.

Standard Trunking Criteria

Observe the following criteria for standard distance trunking:

- There must be a direct connection between participating switches.
- Trunk ports must reside in the same port group.
- Trunk ports must run at the same speed (either 2 Gbit/sec or 4 Gbit/sec).
- Trunk ports must be set to the same ISL mode (L0 is the default). For details on extended ISL modes, see Table 18-1 on page 18-4.
- Trunk ports must be E_Ports or EX_Ports.
- For optimal performance, cable lengths of no more than 30 meters difference are recommended.
- The **switch.interopMode** parameter must be set to 0. See "Configuring Interoperability Mode" on page B-1 for information and procedures related to interoperability mode.
- The port ISL mode must be disabled (using the **portCfgIslMode** command).

EX_Port Trunking

In Fabric OS 5.2.0, you can configure EX_ports to use frame based trunking just like regular E- ports. Ex-port frame based trunking has a “master” trunk link. Should the master link go down, the entire trunk will re-form and it will be taken offline for a short period of time. The EX_port restrictions are the same as E-ports. A trunk of E or Ex-port can be up to 8-ports wide. All the ports must be adjacent to each other using the clearly marked groups on the front of the product.

Fabric Considerations

The ISL Trunking feature is provided with the Fabric OS and can be activated by entering a license key, available from the switch supplier. When the ISL Trunking license is activated (after you have entered the **switchDisable** and **switchEnable** commands), trunking is automatically implemented for any eligible ISLs.

A license must be activated on each switch that participates in trunking.

To use ISL Trunking in the fabric, the fabric must be designed to allow trunking groups to form. To identify the most useful trunking groups, evaluate the traffic patterns before designing or redesigning the fabric. This also applies to the Silkworm 24000 configured with two domains, and the SilkWorm 48000, which does not support two domains.

ISL Trunking can be used to simplify storage area network design and improve performance. When designing the SAN, consider the following recommendations in addition to the standard guidelines for SAN design:

- Evaluate the traffic patterns within the fabric.
- Place trunking-capable switches adjacent to each other.

  This maximizes the number of trunking groups that can form. If you are using a core/edge topology, place trunking-capable switches at the core of the fabric and any switches that are not trunking-capable at the edge of the fabric.
- Activate an ISL Trunking license on each switch that is to participate in a trunking group.
- For optimal performance, no more than 30 meters difference is recommended.
- When connecting two switches with two or more ISLs, ensure that all trunking requirements are met to allow a trunking group to form.
- Determine the optimal number of trunking groups between each set of linked switches, depending on traffic patterns and port availability.

  The goal is to avoid traffic congestion without unnecessarily using ports that could be used to attach other switches or devices. Consider these points:

  - Each physical ISL uses two ports that could otherwise be used to attach node devices or other switches.
  - Trunking groups can be used to resolve ISL oversubscription if the total capability of the trunking group is not exceeded.

- Consider how the addition of a new path will affect existing traffic patterns:
  - A trunking group has the same link cost as the master ISL of the group, regardless of the number of ISLs in the group. This allows slave ISLs to be added or removed without causing data to be rerouted, because the link cost remains constant.
  - The addition of a path that is shorter than existing paths causes traffic to be rerouted through that path.
  - The addition of a path that is longer than existing paths might not be useful because the traffic will choose the shorter paths first.
- Plan for future bandwidth addition to accommodate increased traffic.

  For trunking groups over which traffic is likely to increase as business requirements grow, consider leaving one or two ports in the group available for future nondisruptive addition of bandwidth.
- Consider creating redundant trunking groups where additional ports are available or paths are particularly critical.

  This helps to protect against oversubscription of trunking groups, multiple ISL failures in the same group, and the rare occurrence of an ASIC failure.
- To provide the highest level of reliability, deploy trunking groups in redundant fabrics to further ensure ISL failures do not disrupt business operations.

Initializing Trunking on Ports

After you unlock the ISL Trunking license, you must reinitialize the ports being used for ISLs so that they recognize that trunking is enabled. This procedure only needs to be performed one time.

To reinitialize the ports, you can either disable and then reenable the switch, or disable and then reenable the affected ports.

To disable and reenable the switch

1. Connect to the switch and log in as admin.
2. Enter the **switchDisable** command.
3. Enter the **switchEnable** command.

To disable and reenable ports

1. Connect to the switch and log in as admin.
2. Enter the **portDisable** command. The format is:

   portDisable [*slot*/]*port*

   Slot is the slot number (SilkWorm 24000 directors only) and *port* is the port number of the port you want to disable.
3. Enter the **portEnable** command.The format is:

   portEnable [*slot*/]*port*

   Slot is the slot number (SilkWorm 24000 and 48000 directors only) and *port* is the port number of the port you want to enable.

Monitoring Traffic

To implement ISL Trunking effectively, you must monitor fabric traffic to identify congested paths or to identify frequently dropped links. While monitoring changes in traffic patterns, you can adjust the fabric design accordingly, such as by adding, removing, or reconfiguring ISLs and trunking groups in problem areas.

There are three methods of monitoring fabric traffic:

- *Brocade Advanced Performance Monitoring* monitors traffic flow and allows you to view the impact of different fabric configurations on performance. See "Administering Advanced Performance Monitoring" on page 17-1 for additional information.
- *Brocade Fabric Watch* allows you to monitor traffic flow through specified ports on the switch and send alerts when the traffic exceeds or drops below configurable thresholds. See the *Fabric Watch Administrator's Guide* for additional information.
- Use the **portPerfShow** command as described in the following procedure to record traffic volume for each port in your fabric over time.

To use the portPerfShow command

1. Connect to the switch and log in as admin.
2. Enter the following command:

   ```
   portperfshow [interval]
   ```

   where *interval* is the number of seconds between each data-gathering sample (the default is one sample every second).
3. Record the traffic flow for each port participating in an ISL.
4. Repeat step 1 through step 3 for each switch in the fabric until all ISL traffic flow is captured.

   In a large fabric, it might be necessary to only identify and capture the key ISLs. However, you might want to continue this process throughout the day (or an entire work cycle), to capture varying traffic patterns under different conditions.

The following example shows a switch without trunking, and indicates that ports 0 through 2 are under-utilized and ports 4 and 5 are congested:

```
switch:admin> portperfshow
0       1       2       3       4       5       6       7       Total
-------------------------------------------------------------------------
0       0       0       145m    204m    202m    0       168m    719
0       0       0       145m    206m    208m    0       186m    745
switch:admin>
```

The following example shows traffic flowing through a trunking group (ports 5, 6, and 7). After port 6 fails, traffic is redistributed over the remaining two links in the group, ports 5 and 7:

```
switch:admin> portperfshow
0       1       2       3       4       5       6       7       Total
-------------------------------------------------------------------------
0       0       0       0       0       145m    144m    145m    434
0       0       0       0       0       144m    143m    144m    431
0       0       0       0       0       162m    0       162m    324
0       0       0       0       0       186m    0       186m    372
0       0       0       0       0       193m    0       192m    385
0       0       0       0       0       202m    0       202m    404
0       0       0       0       0       209m    0       209m    418
switch:admin>
```

Enabling and Disabling ISL Trunking

You can enable or disable Brocade ISL Trunking for a single port or for an entire switch.When you execute the commands **portCfgTrunkPort** or **switchCfgTrunk** to update the trunking configuration, the ports for which the configuration applies are disabled and reenabled with the new trunk configuration. As a result, traffic through those ports could be disrupted.

To enable or disable ISL Trunking on one port

1. Connect to the switch and log in as admin.
2. Enter the **portCfgTrunkPort** command. The format is:

```
portcfgtrunkport [slotnumber/]portnumber mode
```

slotnumber	Specifies the number of the slot in which the port blade containing the port is located (this operand only required for switches with slots such as the SilkWorm 24000, and 48000 directors).
portnumber	Specifies the number of the port on which you want to enable or disable trunking.
mode	Enables (1) or disables (0) trunking on the specified port.

The following example enables trunking on slot 1, port 3:

```
switch:admin> portcfgtrunkport 1/3 1
done.
switch:admin>
```

To enable or disable ISL Trunking for all of the ports on a switch

1. Connect to the switch and log in as admin.
2. Enter the **switchCfgTrunk** command. The format is:

```
switchcfgtrunk mode
```

Mode 1 enables and mode 0 disables ISL Trunking for all ports on the switch.

The following example enables trunking all ports in the switch.

```
switch:admin> switchcfgtrunk 1
Committing configuration...done.
switch:admin>
```

Setting Port Speeds

For long-distance ports, if a port is set to autonegotiate port speed, the maximum speed (which is 4 Gbit/sec) is assumed for reserving buffers for the port–this wastes buffers if the port is actually running at 2 Gbit/sec. For long-distance ports, it is best to set the port speed (this applies to SilkWorm 4100 and SilkWorm 48000 only).

You can set the port speed for one port or for an entire switch. Trunked ports must be set to the same speed.

To set the speed for one port

1. Connect to the switch and log in as admin.
2. Enter the **portCfgSpeed** command. The format is:

```
portcfgspeed [slotnumber/]portnumber, speed_level
```

slotnumber	For bladed systems only, specify the slot number of the port to be configured, followed by a slash (/). This operand is only required for switches with slots, such as the SilkWorm 24000 and 48000.
portnumber	Specifies the port number relative to its slot for bladed systems.
speedlevel	Specifies the speed of the link: • 0—Autonegotiating mode. The port automatically configures for the highest speed. • 1—one Gbit/sec mode. Fixes the port at a speed of one Gbit/sec. Changing the speed to one Gbit/sec causes the port to be excluded from the trunk group. • 2—two Gbit/sec mode. Fixes the port at a speed of two Gbit/sec. • 4—four Gbit/sec mode. Fixes the port at a speed of four Gbit/sec. (SilkWorm 200E, 4100 and 48000 only.)

The following example sets the speed for port 3 on slot 2 to two Gbit/sec:

```
switch:admin> portcfgspeed 2/3 2
done.
switch:admin>
```

The following example sets the speed for port 3 on slot 2 to autonegotiate:

```
switch:admin> portcfgspeed s/3 0
done.
switch:admin>
```

To set the speed for all of the ports on the switch

1. Connect to the switch and log in as admin.
2. Enter the **switchCfgSpeed** command. The format is:

```
switchcfgspeed speedlevel
```

speedlevel Specifies the speed of the link:

- 0—Auto-negotiating mode. The port automatically configures for the highest speed.
- 1—one Gbit/sec mode. Fixes the port at a speed of one Gbit/sec. Changing the speed to one Gbit/sec causes the port to be excluded from the trunk group.
- 2—two Gbit/second mode. Fixes the port at a speed of two Gbit/second.
- 4—four Gbit/second mode. Fixes the port at a speed of four Gbit/second. (SilkWorm 200E, 4100 and 48000 switches only.)

The following example sets the speed for all ports on the switch to two Gbit/second:

```
switch:admin> switchcfgspeed 2
Committing configuration...done.
switch:admin>
```

The following example sets the speed for all ports on the switch to autonegotiate:

```
switch:admin> switchcfgspeed 0
Committing configuration...done.
switch:admin>
```

Displaying Trunking Information

The **trunkShow** command offers an efficient means of listing out all the trunks and members of a trunk. You can easily discover the peer ports for disabling a port, disconnecting or adding additional members. Viewing trunks is also useful for when you want to ensure that trunks are formed correctly.

Use the **trunkShow** command to display the following information about ISL Trunking groups:

- Number identifier.
- Port-to-port connections, listed in the following format: *local port number -> remote port number.*
- WWNs of the remote switches.
- Deskew values (the time difference, in nanoseconds divided by 10, for traffic to travel over each ISL as compared to the shortest ISL in the group). The system automatically sets the minimum deskew value of the shortest ISL to 15.
- Master ports.

To display trunking information:

1. Connect to the switch and log in as admin.
2. Enter the **trunkShow** command.

This example shows three trunking groups (1, 2, and 3); ports 1, 4, and 14 are masters:

```
 switch:admin> trunkshow
 1: 1 ->  1   10:00:00:60:69:04:10:83   deskew 16 Master
    0 ->  0   10:00:00:60:69:04:10:83   deskew 15
 2: 4 ->  4   10:00:00:60:69:04:01:94   deskew 16 Master
    5 ->  5   10:00:00:60:69:04:01:94   deskew 15
    7 ->  7   10:00:00:60:69:04:01:94   deskew 17
    6 ->  6   10:00:00:60:69:04:01:94   deskew 16
 3:14 -> 14   10:00:00:60:69:04:10:83   deskew 16 Master
   15 -> 15   10:00:00:60:69:04:10:83   deskew 15
switch:admin>
```

Trunking Over Extended Fabrics

In addition to the criteria listed in "Standard Trunking Criteria" on page 19-2, observe the following criteria for trunking over extended fabrics:

- ISL Trunking over extended fabrics is supported on switches running Fabric OS v3.2.0 (or later) or v4.4.0 (or later).
- Extended Fabrics and ISL Trunking licenses are required on all participating switches.
- The **vc_translation_link_init** parameter must be set the same on all ports in an extended trunk. (For details on this parameter, see page 18-11.)

Trunking Distances

Enhanced trunking support for switches with Bloom ASICs is summarized in Table 19-2.

Table 19-2 Trunking Support for Bloom ASICs

Mode	Distance	Number of 2 Gbit/sec ports
LE	10 km	4 (one 4-port trunk)
L0.5	25 km	3 (one 3-port trunk)
L1	50 km	1 (one 2-port trunk)
L2	100 km	0

Enhanced trunking support for switches with Condor ASICs is summarized in Table 19-3.

Table 19-3 Trunking Support for SilkWorm 4100 and 4900 (Condor ASIC)

Mode	Distance	Number of 2 Gbit/sec ports	Number of 4 Gbit/sec ports
LE	10 km	32 (four 8-port trunks)	32 (four 8-port trunks)
L0.5	25 km	32 (four 8-port trunks)	15 (one 8-port trunk)
L1	50 km	15 (one 2-port trunk)	7 (one 7-port trunk)
L2	100 km	7 (one 7-port trunk)	3 (one 3-port trunk)
LD	200 km	3 (one 3-port trunk)	0
LD	250 km	3 (one 3-port trunk)	0
LD	500 km	0	0

Enhanced trunking support for the FC4-48 port blade in the SilkWorm 48000 is summarized in Table 19-3.

Table 19-4 Trunking Support for FC4-48 Port Blade (Condor ASIC)

Mode	Distance	Number of 2 Gbit/sec ports	Number of 4 Gbit/sec ports
LE	10 km	48 (six 8-port trunks)	48 (six 8-port trunks)
L0.5	25 km	48 (six 8-port trunks)	22 (one 8-port trunk and one 3-port trunk per Condor)
L1	50 km	22 (one 8-port trunk and one 3-port trunk per Condor)	10 (one 5-port trunk per Condor)
L2	100 km	10 (one 5-port trunk per Condor)	4 (one 2-port trunk per Condor)
LD	200 km	4 (one 2-port trunk per Condor)	0
LD	250 km	4 (one 2-port trunk per Condor)	0
LD	500 km	0	0

Troubleshooting Trunking Problems

If you have difficulty with trunking, try the solutions in this section.

Listing Link Characteristics

If a link that is part of an ISL Trunk fails, use the **trunkDebug** command to troubleshoot the problem, as shown in the following procedure:

1. Connect to the switch and log in as admin.
2. Enter the following command:

```
trunkDebug port, port
```

port Specifies the number of a port in an ISL Trunking group.

The **trunkDebug** command displays the possible reason that two ports cannot be trunked, including the following reasons:

- The switch does not support trunking.
- A trunking license is required.
- Trunking is not supported in switch interoperability mode.
- Port trunking is disabled.
- The port is not an E_Port.
- The port is not 2 Gbit/sec or 4 Gbit/sec.
- The port connects to different switches.
- The ports are not same speed, or they are not set to a valid speed.
- The ports are not set to the same long distance mode.
- Local or remote ports are not in same port group.
- The difference in the cable length among trunked links is greater than the allowed difference.

This example shows that port 3 is not configured as an E_Port:

```
switch:admin> trunkdebug 3 5
port 3 is not E port
switch:admin>
```

Recognizing Buffer Underallocation

For SilkWorm 3250, 3850, 3900 and 24000 models, if there is an underallocation or overcommitment of buffers to ports configured for extended trunking, the switches at both ends of the trunk try to disable some ports so that others can operate using the available buffers. (Standard trunks are not affected by buffer allocation.)

This issue of buffer underallocation does not apply to the SilkWorm 4100 and SilkWorm 48000 models.

A port disabled at one end because of buffer underallocation causes all the disabled ports at the other end to become enabled. Some of these enabled ports become disabled due to a lack of buffers, which in turn triggers ports to be enabled once again at the other end. While the system is stabilizing the buffer allocation, it warns that ports are disabled due to lack of buffers, but it does not send a message to the console when buffers are enabled. The system requires a few passes to stabilize the buffer allocation. Ultimately, the number of ports for which buffers are available come up and stabilize. You should wait for stabilization, and then proceed with correcting the buffer allocation situation.

To get out of buffer-limited mode on E_ports or LD_ports:

1. Change the LD/L1/L2/L0.5 port speed to a lower speed (of non-buffer limited ports).
2. Change the LD port's estimated distance to a shorter distance (of non-buffer limited ports).
3. Change LD/L1/L2/L0.5 back to L0 (of non-buffer limited ports).
4. If you are in buffer-limited mode on the LD port, then increase the estimated distance.

These changes are implemented only after you disable (**portDisable**) and enable (**portEnable**) the buffer-limited port (or buffer-limited switch).

Reconfiguring a port to LD from another mode can result in the port being disabled for lack of buffers—this does not apply to the SilkWorm 4100 and SilkWorm 48000 (using FC4-16 and FC4-32 port blades).

If this happens:

- In Fabric OS v4.2.x, reconfigure the disabled LD port back to the original mode.
- In Fabric OS v4.4.0 and later, specify a slightly shorter distance for the *desired_distance* parameter in the **portCfgLongDistance** command.

Chapter
20

Administering Advanced Zoning

This chapter describes the concepts and procedures for using the Brocade Advanced Zoning feature and contains the following topics:

- "About Zoning" on page 20-1
- "Creating and Managing Zone Aliases" on page 20-11
- "Creating and Maintaining Zones" on page 20-13
- "Creating and Modifying Zoning Configurations" on page 20-19
- "Managing Zoning Configurations in a Fabric" on page 20-26
- "Using Zoning to Administer Security" on page 20-28
- "Resolving Zone Conflicts" on page 20-29

About Zoning

Zoning enables you to partition your storage area network (SAN) into logical groups of devices that can access each other. For example, you can partition your SAN into two zones, *winzone* and *unixzone*, so that your Windows servers and storage do not interact with your UNIX servers and storage.

Zones can be configured dynamically. They can vary in size, depending on the number of fabric-connected devices, and devices can belong to more than one zone. Because zone members can access only other members of the same zone, a device not included in a zone is not available to members of that zone.

When using a mixed fabric—that is, a fabric containing two or more switches running different fabric operating systems—you should use the switch with the highest Fabric OS level to perform zoning tasks.

You must install Brocade Advanced Zoning licenses on all the switches in the fabric before attempting to configure zones. If a Zoning license is removed, you must make sure it is reinstalled properly on the affected switch before attempting the **cfgEnable** zoning operation. Failure to follow these steps can cause inconsistency of the zoning configuration on the affected switches should a zoning operation be attempted from a remote switch in the fabric. On the affected switches, an error message indicates that the Zoning license is missing.

You can use zones to logically consolidate equipment for efficiency or to facilitate time-sensitive functions; for example, use zoning to create a temporary zone to back up nonmember devices.

Any zone object connected to the fabric can be included in one or more zones. Zone objects can communicate only with objects within the same zone. For example, consider Figure 20-1, which shows configured zones, Red, Green, and Blue.

- Server 1 can communicate only with the Loop 1 devices.
- Server 2 can communicate only with the RAID and Blue zone devices.
- Server 3 can communicate with the RAID device and the Loop1 device.
- The Loop 2 JBODs are not assigned to a zone; no other zoned fabric device can access them.

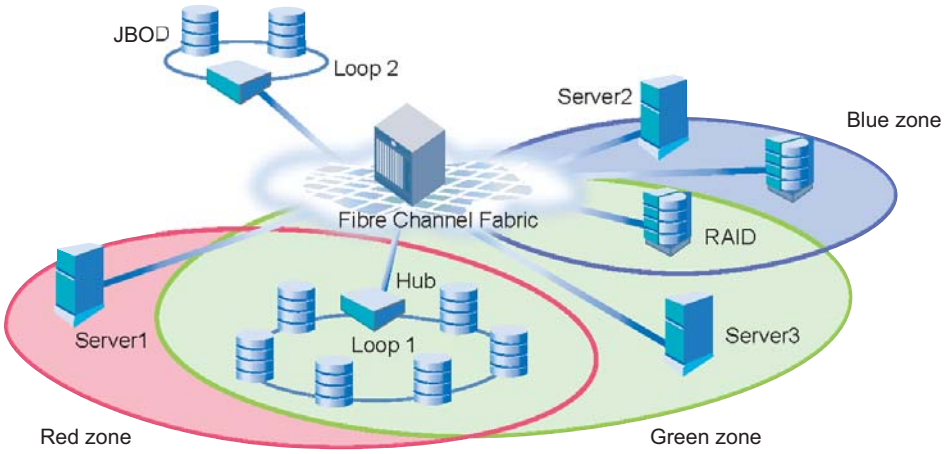

Figure 20-1 Zoning Example

To list the commands associated with zoning, use the **zoneHelp** command.

For detailed information on the zoning commands used in the procedures, see the *Fabric OS Command Reference* or to the online man page for each command.

Before using the procedures, you should become familiar with the zoning concepts described in the following sections.

Zone Types

The following table summarize the types of zoning available.

Table 20-2 Types of Zoning

Zone Type	Description
Storage-based	Storage units typically implement LUN-based zoning, also called *LUN masking*. LUN-based zoning limits access to the LUNs on the storage port to the specific WWN of the server HBA. It is needed in most SANs. It functions during the probe portion of SCSI initialization. The server probes the storage port for a list of available LUNs and their properties. The storage system compares the WWN of the requesting HBA to the defined zone list, and returns the LUNs assigned to the WWN. Other LUNs on the storage port are not made available to the server.
Host-based	Host-based zoning can implement WWN or LUN masking.
Fabric-based	Fabric switches implement fabric-based zoning, in which the zone members are identified by WWN or port location in the fabric. Fabric-based zoning is also called *name server-based* or *soft* zoning. When a device queries the fabric name server, the name server determines the zones in which the device belongs. The server returns information on all members of the zones in the fabric to the device. Devices in the zone are identified by node WWN, port WWN, or (domain, port) of the switch to which the device is connected. The primary approaches to fabric-based zoning are summarized in Table 20-3.

Table 20-3 lists the various approaches you can take when implementing zoning in a fabric.

Table 20-3 Approaches to Fabric-Based Zoning

Zoning Approach	Description
Single HBA	Zoning by single HBA most closely re-creates the original SCSI bus. Each zone created has only one HBA (initiator) in the zone; each of the target devices is added to the zone. Typically, a zone is created for the HBA and the disk storage ports are added. If the HBA also accesses tape devices, a second zone is created with the HBA and associated tape devices in it. In the case of clustered systems, it could be appropriate to have an HBA from each of the cluster members included in the zone; this is equivalent to having a shared SCSI bus between the cluster members and presumes that the clustering software can manage access to the shared devices. In a large fabric, zoning by single HBA requires the creation of possibly hundreds of zones; however, each zone contains only a few members. Zone changes affect the smallest possible number of devices, minimizing the impact of an incorrect zone change. This zoning philosophy is the preferred method.
Application	Zoning by application typically requires zoning multiple, perhaps incompatible, operating systems into the same zones. This method of zoning creates the possibility that a minor server in the application suite could disrupt a major server (such as a Web server disrupting a data warehouse server). Zoning by application can also result in a zone with a large number of members, meaning that more notifications, such as registered state change notifications (RSCNs), or errors, go out to a larger group than necessary.
Operating system	Zoning by operating system has issues similar to zoning by application. In a large site, this type of zone can become very large and complex. When zone changes are made, they typically involve applications rather than a particular server type. If members of different operating system clusters can see storage assigned to another cluster, they might attempt to own the other cluster's storage and compromise the stability of the clusters.
Port allocation	Avoid zoning by port allocation unless the administration team has very rigidly enforced processes for port and device allocation in the fabric. It does, however, provide some positive features. For instance, when a storage port, server HBA, or tape drive is replaced, the change of WWN for the new device is of no consequence. As long as the new device is connected to the original port, it continues to have the same access rights. The ports on the edge switches can be pre-associated to storage ports, and control of the fan-in ratio (the ratio of the input port to output port) can be established. With this pre-assigning technique, the administrative team cannot overload any one storage port by associating too many servers with it.
No fabric zoning	Using no fabric zoning is the least desirable zoning option because it allows devices to have unrestricted access on the fabric. Additionally, any device attached to the fabric, intentionally or maliciously, likewise has unrestricted access to the fabric. This form of zoning should only be utilized in a small and tightly controlled environment, such as when host-based zoning or LUN masking is deployed.

Zone Objects

A *zone object* is any device in a zone, such as the:

- Physical port number or port index on the switch
- Node World Wide Name (N-WWN)
- Port World Wide Name (P-WWN)

Zone objects identified by port number or index number are specified as a pair of decimal numbers in the form *d, index* (*d* is the domain ID of the switch and *index* is the index number on that switch in relation to the port you want to specify).

For example, on SilkWorm 24000 directors, “4, 30” specifies port 14 in slot number 2 (domain ID 4, port index 30). On fixed-port models, “3,13” specifies port 13 in switch domain ID 3.

When a zone object is the physical port number then all devices connected to that port are in the zone.

World Wide Names are specified as 8-byte (16-digit) hexadecimal numbers, separated by colons (:) for example, 10:00:00:90:69:00:00:8a.

When a zone object is the node WWN name, only the specified device is in the zone.

When a zone object is the port WWN name, only the single port is in the zone.

The types of zone objects used to define a zone can be mixed. For example, a zone defined with the zone objects 2,12; 2,14; 10:00:00:80:33:3f:aa:11 contains the devices connected to domain 2, ports 12 and 14, and a device with the WWN (either node name or port name) 10:00:00:80:33:3f:aa:11 that is connected on the fabric.

Zoning Schemes

You can establish a zone by identifying zone objects using one or more of the following *zoning schemes*:

- **Domain, port number level—**All members are specified by *domain ID, port number*, or *domain, index number* pair or aliases, described in “Zone Aliases” on page 20-5.
- **World Wide Name (WWN) level—**All members are specified only by World Wide Name (WWNs) or aliases of WWNs. They can be node or port versions of the WWN.
- **Mixed zoning—**A zone containing members specified by a combination of *domain, port number*, and/or *domain, index number* and WWN.

Zone Aliases

A *zone alias* is a name assigned to a device or a group of devices. By creating an alias, you can assign a familiar name to a device or group multiple devices into a single name. This simplifies cumbersome data entry and allows an intuitive naming structure (such as using “NT_Hosts” to define all NT hosts in the fabric).

Zone aliases also simplify repetitive entry of zone objects such as port numbers or a WWN. For example, you can use the name “Eng” as an alias for “10:00:00:80:33:3f:aa:11”.

Name zones for the initiator they contain can also be useful. For example, if you use the alias SRV_MAILSERVER_SLT5 to designate a mail server in PCI slot 5, then the alias for the associated zone is ZNE_MAILSERVER_SLT5. This clearly identifies the server host bus adapter (HBA) associated with the zone.

Zone configuration naming is flexible. One configuration should be named PROD_*fabricname*, where *fabricname* is the name that the fabric has been designated. The purpose of the PROD configuration is to easily identify the configuration that can be implemented and provide the most generic services. If other configurations are used for specialized purposes, names such as “BACKUP_A,” “RECOVERY_2,” and “TEST_18jun02” can be used.

Zone Configurations

A *zone configuration* is a group of one or more zones. A zone can be included in more than one zone configuration. When a zone configuration is in effect, all zones that are members of that configuration are in effect.

The different types of zone configurations are:

- **Defined Configuration—**The complete set of all zone objects defined in the fabric.
- **Effective Configuration—**A single zone configuration that is currently in effect. The effective configuration is built when you enable a specified zone configuration.
- **Saved Configuration—**A copy of the defined configuration plus the name of the effective configuration, which is saved in flash memory by the **cfgSave** command. (You can also use the **configUpload** command to provide a backup of the zoning configuration and the **configDownload** command to restore the zoning configuration.) There might be differences between the saved configuration and the defined configuration if you have modified any of the zone definitions and have not saved the configuration.
- **Disabled Configuration—**The effective configuration is removed from flash memory.

On power-up, the switch automatically reloads the saved configuration. If a configuration was active when it was saved, the same configuration is reinstated on the local switch with an autorun of the **cfgEnable** command.

Zoning Enforcement

Zoning enforcement describes a set of predefined rules that the switch uses to determine where to send incoming data. There are two methods of enforcement which are hardware and software-enforced zoning.

Hardware-Enforced Zoning

Hardware-enforced zoning is specified without using the mixed zoning scheme (mixed zones contain domains, ports and WWNs as zone members). The exact methodology varies on different switch models.

Hardware-enforced zoning

- Is also called “hard zoning.”
- Prevents a host from accessing a device it is not authorized to access.

- Checks each frame before it is delivered to a zone member and discards it if there is a zone mismatch. When hardware-enforced zoning is active, the Brocade switch monitors the communications and blocks any frames that do not comply with the effective zone configuration. The switch performs this blocking at the transmit side of the port on which the destination device is located.
- Is enforced at the ASIC level. Each ASIC maintains a list of source port IDs that have permission to access any of the ports on that ASIC.
- Is available on 1 Gbit/sec, 2 Gbit/sec, and 4 Gbit/sec platforms.
- Prevents hosts from discovering unauthorized target devices.
- Ensures that the name server does not return any information to an unauthorized initiator in response to a name server query.
- Is exclusively enforced through selective information presented to end nodes through the fabric Simple Name Server (SNS). When an initiator queries the name server for accessible devices in the fabric, the name server returns only those devices that are in the same zone as the initiator. Devices that are not part of the zone are not returned as accessible devices.

Fabric OS uses hardware-enforced zoning (on a per-zone basis) whenever the fabric membership or zone configuration changes.

Software-Enforced Zoning

Zoning enables you to restrict access to devices in a fabric. Software-enforced zoning prevents hosts from discovering unauthorized target devices, while hardware-enforced zoning prevents a host from accessing a device it is not authorized to access.

Software-enforced zoning:

- Is also called “soft zoning,” “Name Server zoning,” “fabric-based zoning,” “session-based zoning,” or “hardware-assisted zoning.”
- Is available on 1-, 2-, and 4-Gbit/sec platforms.
- Prevents hosts from discovering unauthorized target devices.
- Ensures that the name server does not return any information to an unauthorized initiator in response to a name server query.
- Does not prohibit access to the device. If an initiator has knowledge of the network address of a target device, it does not need to query the Name Server to access it, which could lead to undesired access to a target device by unauthorized hosts.
- Is exclusively enforced through selective information presented to end nodes through the fabric Simple Name Server (SNS). When an initiator queries the name server for accessible devices in the fabric, the name server returns only those devices that are in the same zone as the initiator. Devices that are not part of the zone are not returned as accessible devices.

Table 20-4 shows the various Brocade switch models, the hardware zoning methodology for each, and tips for best usage.

Table 20-4 Enforcing Hardware Zoning

Fabric Type	Methodology	Best Practice
SilkWorm 2000-series	Enables hardware-enforced zoning only on domain, port zones; WWN or mixed zones are not hardware-enforced. Any domain, port zone that overlaps a mixed or WWN zone is not hardware-enforced. An overlap occurs when a member specified by WWN is connected to a port in a domain, port zone. The domain, port zone loses its hardware enforcement even though a review of the zone configuration does not indicate it.	Use *domain, port* identifiers. Do not identify a zone member by its WWN.
SilkWorm 3000-series, SilkWorm 200E, 4100, 4900, 7500, 24000, and 48000 models	Enable hardware-enforced zoning on domain, port zones, and WWN zones. Overlap of similar zone types does not result in the loss of hardware enforcement. Overlap with other zone type results in the loss of hardware enforcement. As in the SilkWorm 2000-series switches, connecting a device specified by WWN into a port specified in a domain, port zone results in loss of the hardware enforcement in both zones. **For the SilkWorm 48000 with an FC4-48 port blade:** If ports 16 through 47 on the FC4-48 port blade use *domain, port* identifiers, then session-based zoning is enforced on these ports.	Use either WWN or *domain, port* identifiers.
Mixed switches	Enable hardware-enforced zoning according to each switch type. Use the **portZoneShow** command to find the zone type to which a device is attached.	Use *domain, port* identifiers. You can use WWN identifiers if you place disk and tape targets on SilkWorm 3000-series, and SilkWorm 24000 models, and do not use *domain, port* identifiers.

Figure 20-5 shows a fabric with four nonoverlapping hardware-enforced zones.

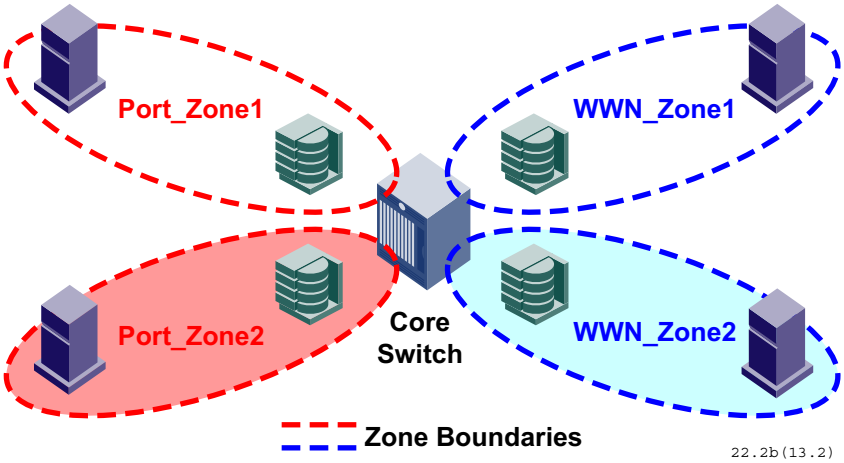

Figure 20-5 Hardware-Enforced Nonoverlapping Zones

Figure 20-6 shows the same fabric components zoned in an overlapping fashion.

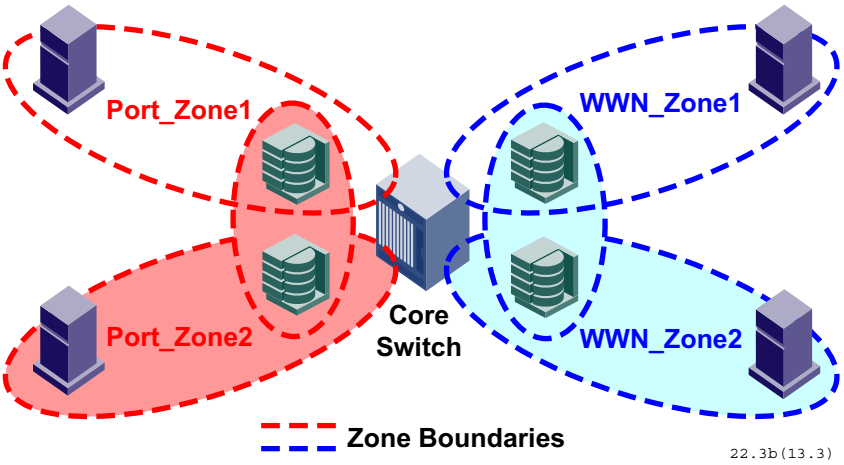

Figure 20-6 Hardware-Enforced Overlapping Zones

Any zone using a mixed zoning scheme on the Brocade 2 Gbit/sec platform rely on Name Server authentication as well as hardware-assisted (ASIC) authentication, which ensures that any PLOGI/ ADISC/PDISC/ACC from an unauthorized device attempting to access a device it is not zoned with is rejected. Brocade 2 Gbit/sec switches always deploy the hardware assist, in any zone configuration (see Figure 20-7 and Figure 20-8).

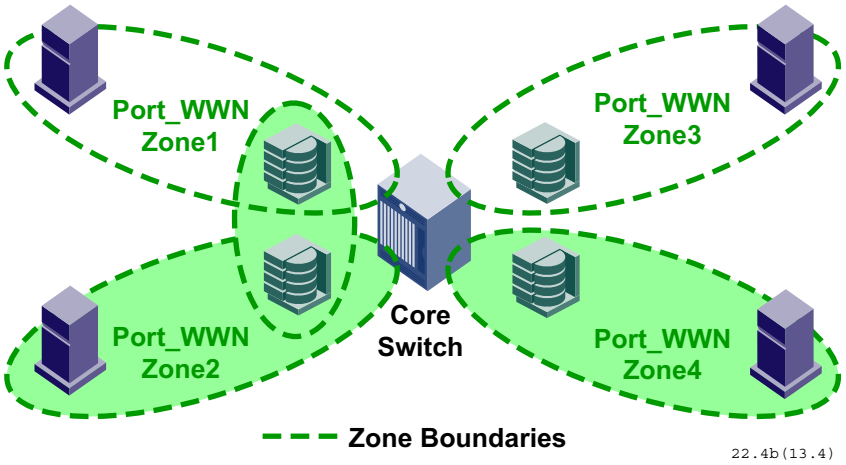

Figure 20-7 Zoning with Hardware Assist (Mixed-Port and WWN Zones)

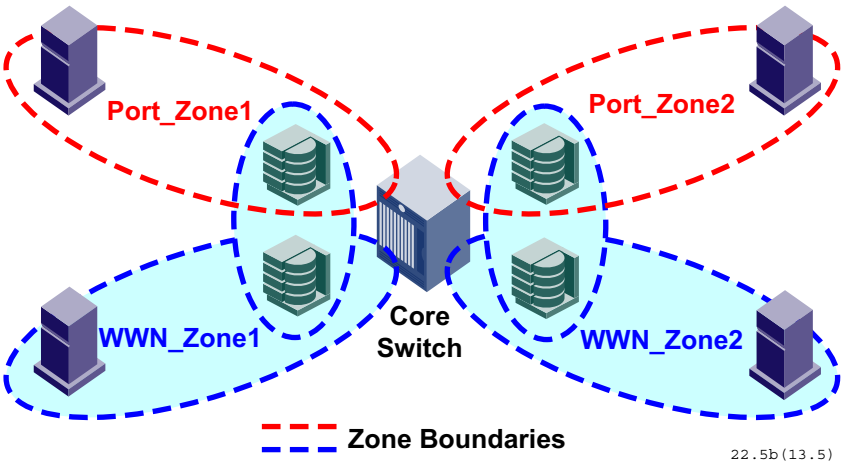

Figure 20-8 Session-Based Hard Zoning

In Figure 20-8, only the overlapping ports are software-enforced with hardware assist.

Rules for Configuring Zones

Observe the following rules when configuring zones:

- If security is a priority, you should use hard zoning.
- Using aliases is optional with zoning, and requires structure when defining zones. However, aliases aid administrators of a zoned fabric to understand the structure and context.
- Evaluate the security requirements of the fabric. If additional security is required, add Brocade Secure Fabric OS into the fabric.
- If the fabric includes a Brocade switch and you support a third-party switch product, they are only able to use WWN zoning; other types of zoning, including QuickLoop, are not supported.

- QuickLoop Fabric Assist: Evaluate whether the fabric uses QuickLoop Fabric Assist (QLFA). If you are running Brocade Fabric OS v4.x, consider that Brocade Fabric OS v4.x cannot have a Fabric Assist host directly connected to it. However, targets on a Brocade Fabric OS v4.x switch can still be part of a Fabric Assist zone if a Fabric Assist host is connected to a non-v4.x switch.
- Zone changes: Zone changes in a production fabric can cause a disruption of I/O when an RSCN is generated because of the zone change and the HBA is unable to process the RSCN fast enough. Although RSCNs are a normal part of a functioning SAN, the pause in I/O might not be acceptable. For these reasons, you should perform zone changes only when the resulting behavior is predictable and acceptable. Changing HBA drivers can shorten the response time in relation to the RSCN.
- Final verification: After changing or enabling a zone configuration, confirm that the nodes and storage can identify and access one another. Depending on the platform, you might need to reboot one or more nodes in the fabric with the new changes.

The zone configuration is managed on a fabric basis. Zoning can be implemented and administered from any switch in the fabric (it is best to use a newer switch such as the SilkWorm 4100, 4900, or a director) that has an Advanced Zoning license enabled. When a change in the configuration is saved, enabled, or disabled per the transactional model, it is automatically (by closing the transaction) distributed to all switches in the fabric, preventing a single point of failure for zone information.

Note

Zoning commands make changes that affect the entire fabric. When executing fabric-level configuration tasks, allow time for the changes to propagate across the fabric before executing any subsequent commands. For a large fabric, you might want to wait several minutes between commands.

Creating and Managing Zone Aliases

A zone alias is a logical group of ports or WWNs. You can simplify the process of creating zones by first specifying aliases, which eliminates the need for long lists of individual zone member names.

If you are creating a new alias using **aliCreate w**, "**1,1**", and a user in another telnet session executes **cfgEnable** (or **cfgDisable**, or **cfgSave**), the other user's transaction will abort your transaction and you will receive an error message. Creating a new alias while there is a zone merge taking place might also abort your transaction. For more details about zone merging and zone merge conflicts, see "Adding a New Switch or Fabric" on page 20-26.

To create an alias

1. Connect to the switch and log in.
2. Enter the **aliCreate** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

```
switch:admin> alicreate "array1", "2,32; 2,33; 2,34; 4,4"
switch:admin> alicreate "array2", "21:00:00:20:37:0c:66:23; 4,3"
switch:admin> alicreate "loop1", "4,6"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

To add members to an alias

1. Connect to the switch and log in.
2. Enter the **aliAdd** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

```
switch:admin> aliadd "array1", "1,2"
switch:admin> aliadd "array2", "21:00:00:20:37:0c:72:51"
switch:admin> aliadd "loop1", "5,6"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

To remove members from an alias

1. Connect to the switch and log in.
2. Enter the **aliRemove** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

```
switch:admin> aliremove "array1", "1,2"
switch:admin> aliremove "array2", "21:00:00:20:37:0c:72:51"
switch:admin> aliremove "loop1", "4,6"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Note

For Fabric OS versions earlier than v4.4.0, when using the **aliRemove** command, the order in which the members appear in the list is critical. For more information on this command, see the *Fabric OS Command Reference*.

To delete an alias

1. Connect to the switch and log in.
2. Enter the **aliDelete** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

```
switch:admin> alidelete "array1"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

To view an alias in the defined configuration

1. Connect to the switch and log in.
2. Enter the **aliShow** command.

   The following example shows all zone aliases beginning with “arr”.

```
switch:admin> alishow “arr*”
alias: array1  21:00:00:20:37:0c:76:8c
alias: array2  21:00:00:20:37:0c:66:23
```

If no parameters are specified, the entire zone database (both the defined and effective configuration) is displayed.

Creating and Maintaining Zones

Before executing **cfgDisable**, **cfgEnable**, or **cfgSave**, execute the **rcsDisabled** command to check whether or not your fabric has Reliable Commit Service (RCS) enabled (*rcsDisabled=0*). If RCS is disabled (*rcsDisabled=1*), check for older switches in the fabric. After the older switches are upgraded, RCS is enabled by default.

RCS is available on all switches running Fabric OS versions 4.1 and later. RCS guarantees that either *all* or *none* of the switches receive the new zone configuration. You should use RCS to secure a reliable propagation of the latest zone configuration. If you use non-RCS mode, you must log in to every switch to monitor the status of the zone configuration.

To create a zone

1. Connect to the switch and log in.
2. Enter the **zoneCreate** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

```
switch:admin> zonecreate “greenzone”, “2,32; 2,33; 2,34; 4,4”
switch:admin> zonecreate “redzone”, “21:00:00:20:37:0c:66:23; 4,3”
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

To add devices (members) to a zone

1. Connect to the switch and log in.
2. Enter the **zoneAdd** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

```
switch:admin> zoneadd "greenzone", "1,2"
switch:admin> zoneadd "redzone", "21:00:00:20:37:0c:72:51"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

To remove devices (members) from a zone

1. Connect to the switch and log in.
2. Enter the **zoneRemove** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

```
switch:admin> zoneremove "greenzone", "1,2"
switch:admin> zoneremove "redzone", "21:00:00:20:37:0c:72:51"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

To delete a zone

1. Connect to the switch and log in.
2. Enter the **zoneDelete** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

```
switch:admin> zonedelete "redzone"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

To view a zone in the defined configuration

1. Connect to the switch and log in.
2. Enter the **zoneShow** command.

   The following example shows all zones beginning with A, B, or C:

```
switch:admin> zoneshow "[A-C]*"
zone: Blue_zone 1,1; array1; 1,2; array2
zone: Bobs_zone 4,5; 4,6; 4,7; 4,8; 4,9
```

If no parameters are specified, the entire zone database (both the defined and effective configuration) is displayed.

Activating Default Zones

Typically, when you issue the **cfgDisable** command in a large fabric with thousands of devices, the name server indicates to all hosts that they can communicate with each other. In fact, each host can receive an enormous list of PIDs, and ultimately cause other hosts to run out of memory or crash. To ensure that all devices in a fabric do not see each other during a **cfgDisable** operation, you can activate a *default zone*.

A default zone includes all devices in a fabric that are not members of an active zone set. Depending upon the configuration, default zones are activated whenever you enter the **cfgDisable** command. Default zones have the following properties:

- Any Nx_Port that is not a member of any zone in the active zone set is a member of the default zone.
- The fabric does not permit members of the default zone to interact with members of any other zone.
- The fabric might not allow members of the default zone to interact with other members of the default zone.

When you activate default zoning a new zone alias is created in the defined zone configuration and is propagated to the fabric when you type the **cfgSave** command. Then, when you issue the **cfgDisable** command, if the zone alias exists, zoning actually interprets the **cfgDisable** command as a **cfgEnable** command for the default zone. The default zone applies to the entire fabric (regardless of switch model).

To activate a default zone

1. Connect to the switch and log in.
2. Enter the **cfgActvShow** command to view the current zone configuration.
3. If no zone configuration transactions are in progress, then enter the **defZone --noaccess** command, which prevents all Nx_Ports from communicating with each other.
4. Enter either the **cfgSave**, **cfgEnable**, or **cfgDisable** command to commit the change and distribute it to the fabric. The change will not be committed and distributed across the fabric if you do not enter one of these commands.

After setting the default zone access mode to *noaccess*, subsequent use of the **cfgDisable** command performs the equivalent of the **cfgEnable "d_efault_Cfg"** command.

To view the current default zone access mode

1. Connect to the switch and log in.
2. Enter the **defZone --show** command.

Note

If you performed a firmware download of an older release, then the current default zone access state will appear as it did prior to the download. For example, if the d_efault_Cfg was in effect before the download, it will remain in effect afterward.

See the *Fabric OS Command Reference* for additional information on the **defZone** command.

Merging Zones

Table 20-9 presents zoning database size limitations for various Fabric OS release versions. The maximum size of a zone database is the upper limit for the defined configuration, and it is determined by the amount of flash memory available for storing the defined configuration.

Table 20-9 Zoning Database Limitations

Fabric OS Version	Maximum Database Size (KB))
2.4.0	64
2.5.0	64
2.6.0	96
3.0.0	128
3.1.0	96
3.2.0	256
4.0.0, 4.1.0, 4.2.0	128
4.4.0	256
5.0.1	256
5.1.0	256
5.2.0	1024 (see the following note for restrictions)

Before linking two switches together, it is important that you know the zone database limit of adjacent switches. For example, when switches running Fabric OS v3.2, v4.4.0, or v5.1.0 discover that the zone merge database is larger than its pre-determined zone database size limit, they issue a reject notification before symmetrically segmenting their own ends of the ISL, thereby preventing the new switch from joining the fabric.

Note

Fabric OS 5.2.0 supports a zoning database size of 1 MB. To support a 1 MB zoning database, the entire fabric must be upgraded to v5.2.0. If any pre-5.2.0 switches are part of the fabric, the maximum zone database size is 256 KB. Switches not upgraded to Fabric OS v5.2.0 are segmented out of the fabric, if the zoning database size exceeds 256 KB (use the **cfgSize** command to display the zoning database size).

Symmetrical segmentation occurs when both ends of an ISL are shut down. Subsequently, no frames are exchanged between those two switches.

Asymmetrical segmentation not only prevents frames from being exchanged between switches, but also causes routing inconsistencies.

The best way to avoid either type of segmentation is to know the zone database size limit of adjacent switches. The following tables provide the expected behavior based on different database sizes after a zone merge is specified.

Table 20-10 Resulting Database Size: 0 to 96K

Receiver / Initiator	Fabric OS v2.6	Fabric OS v3.1	Fabric OS v3.2	Fabric OS v4.0/ v4.1/ v4.2	Fabric OS v4.3/ v4.4.0	Fabric OS v5.0.0/ v5.0.1/ v5.1.0	Fabric OS v5.2.0	Fibre Channel Router	XPath v7.3
Fabric OS v2.6/v3.1	Join	Join	Join	Join	Join	Join	Join	Join	Join
Fabric OS v3.2	Join	Join	Join	Join	Join	Join	Join	Join	Join
Fabric OS v4.0/ v4.1/v4.2	Join	Join	Join	Join	Join	Join	Join	Join	Join
Fabric OS v4.4.0	Join	Join	Join	Join	Join	Join	Join	Join	Join
Fabric OS v5.0.0/ v5.0.1/v5.1.0	Join	Join	Join	Join	Join	Join	Join	Join	Join
Fabric OS v5.2.0	Join	Join	Join	Join	Join	Join	Join	Join	Join
Fibre Channel Router	Join	Join	Join	Join	Join	Join	Join	Join	Join
XPath v7.3	Join	Join	Join	Join	Join	Join	Join	Join	Join

Table 20-11 Resulting Database Size: 96K to 128K

Receiver / Initiator	Fabric OS v2.6	Fabric OS v3.1	Fabric OS v3.2	Fabric OS v4.0/ v4.1/ v4.2	Fabric OS v4.3/ v4.4.0	Fabric OS v5.0.0/ v5.0.1/ 5.1.0	Fabric OS v5.2.0	Fibre Channel Router	XPath v7.3
Fabric OS v2.6/ v3.1	Segment	Segment	Segment	Segment	Segment	Segment	Join	Join	Segment
Fabric OS v3.2	Segment	Segment	Join	Join	Join	Join	Join	Join	Join
Fabric OS v4.0/ v4.1/v4.2	Segment	Segment	Segment	Join	Join	Join	Join	Join	Join
Fabric OS v4.4.0	Segment	Segment	Join	Join	Join	Join	Join	Join	Join
Fabric OS v5.0.0/ v5.0.1	Segment	Segment	Join	Join	Join	Join	Join	Join	Join
Fabric OS v5.2.0	Segment	Segment	Join	Join	Join	Join	Join		
Fibre Channel Router	Segment	Segment	Join	Join	Join	Join	Join	Join	Join
XPath v7.3	Segment	Segment	Segment	Join	Join	Join	Join	Join	Join

Table 20-12 Resulting Database Size: 128K to 256K

Receiver / Initiator	Fabric OS v2.6	Fabric OS v3.1	Fabric OS v3.2	Fabric OS v4.0/ v4.1/ v4.2	Fabric OS v4.3/ v4.4.0	Fabric OS v5.0.0/ v5.0.1/ 5.1.0	Fabric OS v5.2.0	Fibre Channel Router	XPath v7.3
Fabric OS v2.6/ v3.1	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Join	Segment
Fabric OS v3.2	Segment	Segment	Join	Segment	Join	Join	Join	Join	Segment
Fabric OS v4.0/ v4.1/v4.2	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment
Fabric OS v4.4.0	Segment	Segment	Join	Segment	Join	Join	Join	Join	Segment
Fabric OS v5.0.0/ v5.0.1	Segment	Segment	Join	Segment	Join	Join	Join	Join	Segment
Fabric OS v5.2.0	Segment	Segment	Join	Join	Join	Join	Join	Join	Join
Fibre Channel Router	Segment	Segment	Join	Segment	Join	Join	Join	Join	Segment
XPath v7.3	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment

Table 20-13 Resulting Database Size: 256K to 1M

Receiver / Initiator	Fabric OS v2.6	Fabric OS v3.1	Fabric OS v3.2	Fabric OS v4.0/ v4.1/ v4.2	Fabric OS v4.3/ v4.4.0	Fabric OS v5.0.0/ v5.0.1/ v5.1.0	Fabric OS v5.2.0	Fibre Channel Router	XPath v7.3
Fabric OS v2.6/ v3.1	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment
Fabric OS v3.2	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment
Fabric OS v4.0/ v4.1/v4.2	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment
Fabric OS v4.4.0	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment
Fabric OS v5.0.0/v5.0.1/ v5.1.0	Segment	Segment	Segment	Asymmetrical Segment	Segment	Join	Segment	Segment	Segment
Fabric OS v5.2.0	Segment	Segment	Segment	Segment	Segment	Segment	Join	Join	Segment
Fibre Channel Router	Segment	Segment	Segment	Segment	Segment	Join	Join	Join	Segment
XPath v7.3	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment

Creating and Modifying Zoning Configurations

You can store a number of zones in a zoning configuration database. The maximum number of items that can be stored in the zoning configuration database depends on the following criteria:

- Number of switches in the fabric.
- Whether or not interoperability mode is enabled.
- Number of bytes per item. The number of bytes required for an item depends on the specifics of the fabric, but cannot exceed 64 bytes per item.

When enabling a new zone configuration, you must ensure that the size of the configuration does not exceed the minimum size supported by all switches in the fabric. This is particularly important if and when you downgrade to a Fabric OS version that supports a smaller zone database than the current Fabric OS. In this scenario, the zone database in the current Fabric OS would have to be changed to the smaller zone database before the downgrade.

You can use the **cfgSize** command to check both the maximum available size and the currently saved size on all switches. See the *Fabric OS Command Reference* for details on the **cfgSize** command. If you believe you are approaching the maximum, you can save a partially completed zoning configuration and use the **cfgSize** command to determine the remaining space.

For important considerations for managing zoning in a fabric, and more details about the maximum zone database size for each version of the Fabric OS, see "Maintaining Zone Objects" on page 20-23.

To create a zoning configuration

1. Connect to the switch and log in as admin.
2. Enter the **cfgCreate** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

```
switch:admin> cfgcreate "NEW_cfg", "redzone; bluezone; greenzone"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

To add zones (members) to a zoning configuration

1. Connect to the switch and log in.
2. Enter the **cfgAdd** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

```
switch:admin> cfgadd "newcfg", "bluezone"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

To remove zones (members) from a zone configuration

1. Connect to the switch and log in.
2. Enter the **cfgRemove** command.
3. Enter the **cfgSave** command to save the change to the defined configuration.

```
switch:admin> cfgremove "NEW_cfg", "redzone"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

To delete a zone configuration

1. Connect to the switch and log in.
2. Enter the **cfgDelete** command:
3. Enter the **cfgSave** command to save the change to the defined configuration.

```
switch:admin> cfgdelete "testcfg"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

To clear changes to a configuration

1. Use the **cfgTransAbort** command.

   When this command is executed, all changes since the last save operation (performed with the **cfgSave** command) are cleared.

   In the following example, assume that the removal of a member from **zone1** was done in error:

```
switch:admin> zoneremove "zone1","3,5"
switch:admin> cfgtransabort
```

To view all zone configuration information

If you do not specify an operand when executing the **cfgShow** command to view zone configurations, then all zone configuration information (both defined and effective) displays. If there is an outstanding transaction, then the newly edited zone configuration that has not yet been saved is displayed. If there are no outstanding transactions, then the committed zone configuration displays.

1. Connect to the switch and log in as admin.
2. Enter the **cfgShow** command with no operands.

```
switch:admin> cfgshow
Defined configuration:
  cfg:   USA1    Blue_zone
  cfg:   USA_cfg Red_zone; Blue_zone
  zone:  Blue_zone
     1,1; array1; 1,2; array2
  zone:  Red_zone
     1,0; loop1
  alias: array1  21:00:00:20:37:0c:76:8c; 21:00:00:20:37:0c:71:02
  alias: array2  21:00:00:20:37:0c:76:22; 21:00:00:20:37:0c:76:28
  alias: loop1   21:00:00:20:37:0c:76:85; 21:00:00:20:37:0c:71:df

Effective configuration:
  cfg:   USA_cfg
  zone:  Blue_zone
      1,1
      21:00:00:20:37:0c:76:8c
      21:00:00:20:37:0c:71:02
      1,2
      21:00:00:20:37:0c:76:22
      21:00:00:20:37:0c:76:28
   zone:  Red_zone
      1,0
      21:00:00:20:37:0c:76:85
      21:00:00:20:37:0c:71:df
```

To view selected zone configuration information

1. Connect to the switch and log in.
2. Enter the **cfgShow** command and specify a pattern.

   For example, to display all zone configurations that start with “Test”:

```
switch:admin> cfgshow "Test*"
cfg:   Test1 Blue_zone
cfg:   Test_cfg Red_zone; Blue_zone
```

To view a configuration in the effective zone database

1. Connect to the switch and log in.
2. Enter the **cfgActvShow** command.

```
switch:admin> cfgactvshow
Effective configuration:
  cfg:   NEW_cfg
  zone:  Blue_zone
      1,1
      21:00:00:20:37:0c:76:8c
      21:00:00:20:37:0c:71:02
      1,2
      21:00:00:20:37:0c:76:22
      21:00:00:20:37:0c:76:28
   zone:  Red_zone
      1,0
      21:00:00:20:37:0c:76:85
      21:00:00:20:37:0c:71:df
```

Maintaining Zone Objects

While you can use the **cfgDelete** command to delete a zone configuration, there is a quicker and easier way to perform the same task via the zone object commands (**zoneObjectExpunge**, **zoneObjectCopy**, and **zoneObjectRename**). You can also copy and rename zone objects. When you copy a zone object, the resulting object has the same type as the original. Deleting a zone object also removes the object from any member lists of other objects. You can rename objects for all zone object types.

To copy a zone object

1. Connect to the switch and log in.
2. Enter the **cfgShow** command to view the zone configuration objects you want to copy.

   For example, to display all zone configuration objects that start with “Test”:

```
switch:admin> cfgshow "Test*"
cfg:   Test1 Blue_zone
cfg:   Test_cfg Red_zone; Blue_zone
```

3. Enter the **zone --copy** command, specifying the zone configuration objects you want to copy, along with the new object name. Note that zone configuration names are case sensitive; blank spaces are ignored and it works in any Admin Domain other than AD255.

```
switch:admin> zone --copy "Test1", “US_Test1"
```

4. Enter the **cfgShow** command to verify the new zone object is present.

```
switch:admin> cfgshow "Test*"
cfg:   Test1 Blue_zone
cfg:   Test_cfg Red_zone; Blue_zone
switch:admin> cfgShow “US_Test1”
cfg:   Test1 Blue_zone
cfg:   Test_cfg Red_zone; Blue_zone
```

5. If you want the change preserved when the switch reboots, save it to nonvolatile (also known as “flash”) memory by entering the **cfgSave** command.
6. For the change to become effective, enable the appropriate zone configuration using the **cfgEnable** command.

To delete a zone object

1. Connect to the switch and log in.
2. Enter the **cfgShow** command to view the zone configuration objects you want to delete.

```
switch:admin> cfgShow
Defined configuration:
cfg: USA_cfg Red_zone; White_zone; Blue_zone
zone: Blue_zone
     1,1; array1; 1,2; array2
zone: Red_zone
     1,0; loop1
zone: White_zone
     1,3; 1,4
alias: array1 21:00:00:20:37:0c:76:8c; 21:00:00:20:37:0c:71:02
alias: array2 21:00:00:20:37:0c:76:22; 21:00:00:20:37:0c:76:28
alias: loop1 21:00:00:20:37:0c:76:85; 21:00:00:20:37:0c:71:df

Effective configuration:
cfg: USA_cfg
zone: Blue_zone
1,1
21:00:00:20:37:0c:76:8c
21:00:00:20:37:0c:71:02
1,2
21:00:00:20:37:0c:76:22
21:00:00:20:37:0c:76:28
zone: Red_zone
1,0
21:00:00:20:37:0c:76:85
21:00:00:20:37:0c:71:df
```

3. Enter the **zone --expunge** command to delete the zone object. Note that zone configuration names are case sensitive; blank spaces are ignored and it works in any Admin Domain other than AD255.

```
switch:admin> zone --expunge "White_zone"
```

4. Enter the **cfgShow** command to verify the deleted zone object is no longer present.
5. If you want the change preserved when the switch reboots, save it to nonvolatile (also known as “flash”) memory by entering the **cfgSave** command.
6. For the change to become effective, enable the appropriate zone configuration using the **cfgEnable** command.

To rename a zone object

1. Connect to the switch and log in.
2. Enter the **cfgShow** command to view the zone configuration objects you want to rename.

```
switch:admin> cfgShow
Defined configuration:
cfg: USA_cfg Red_zone; White_zone; Blue_zone
zone: Blue_zone
     1,1; array1; 1,2; array2
zone: Red_zone
     1,0; loop1
zone: White_zone
     1,3; 1,4
alias: array1 21:00:00:20:37:0c:76:8c; 21:00:00:20:37:0c:71:02
alias: array2 21:00:00:20:37:0c:76:22; 21:00:00:20:37:0c:76:28
alias: loop1 21:00:00:20:37:0c:76:85; 21:00:00:20:37:0c:71:df
```

3. Enter the **zoneObjectRename** command to rename zone configuration objects. Note that zone configuration names are case sensitive; blank spaces are ignored and it works in any Admin Domain other than AD255.
4. Enter the **cfgShow** command to verify the renamed zone object is present.

```
switch:admin> zoneObjectRename "White_zone", “Red_zone”
```

5. If you want the change preserved when the switch reboots, save it to nonvolatile (also known as “flash”) memory by entering the **cfgSave** command.
6. For the change to become effective, enable the appropriate zone configuration using the **cfgEnable** command.

To validate a zone

1. Connect to the switch and log in.
2. Enter the **cfgShow** command to view the zone configuration objects you want to validate.

```
switch:admin> cfgShow
Defined configuration:
cfg: USA_cfg Red_zone; White_zone; Blue_zone
zone: Blue_zone
     1,1; array1; 1,2; array2
zone: Red_zone
     1,0; loop1
zone: White_zone
     1,3; 1,4
alias: array1 21:00:00:20:37:0c:76:8c; 21:00:00:20:37:0c:71:02
alias: array2 21:00:00:20:37:0c:76:22; 21:00:00:20:37:0c:76:28
alias: loop1 21:00:00:20:37:0c:76:85; 21:00:00:20:37:0c:71:df
```

3. Enter the **zone --validate** command to list all zone members that are not part of the current zone enforcement table. Note that zone configuration names are case sensitive; blank spaces are ignored.

```
switch:admin> zone --validdate "White_zone”
```

4. To validate all zones in the zone database in the defined configuration.

```
switch:admin> sw5:root> zone --validate -m 1
      Defined configuration:
       cfg:   cfg1    zone1
       cfg:   cfg2    zone1; zone2
       zone:  zone1   1,1; ali1
       zone:  zone2   1,1; ali2
       alias: ali1    10:00:00:05:1e:35:81:7f*; 10:00:00:05:1e:35:81:7d*
       alias: ali2    10:00:00:05:1e:35:81:09*; 10:00:00:05:1e:35:81:88*

      ----------------------------------
      ~ - Invalid configuration
      * - Member does not exist
```

   The mode flag -m can be use to specify the zone database location. Supported mode flag values are:

   - 0 - zone database from the current transaction buffer
   - 1 - zone database stored from the persistent storage
   - 2 - currently effective zone database.

   If no mode options are given, the validated output of all the three buffers would be shown.

   If -f option is specified, all the zone members which are not enforceable would be expunged in the transaction buffer. This pruning operation would always happen on the transaction and defined buffers. You cannot specify a mode option or specify a zone object as an argument with -f option. This mode flag should be used after the zone has been validated.

For more details about the **zone**, **cfgShow**, **cfgEnable**, and **cfgSave** commands, see the *Fabric OS Command Reference*.

Managing Zoning Configurations in a Fabric

You can add, delete, or remove individual elements in an existing zone configuration to create an appropriate configuration for your SAN environment. After the changes have been made, save the configuration to ensure the configuration is permanently saved in the switch and that the configuration is replicated throughout the fabric.

The switch configuration file can also be uploaded to the host for archiving and it can be downloaded from the host to a switch in the fabric. See “Backing Up a Configuration” on page 6-2, “Restoring a Configuration” on page 6-5, or the **configUpload** and **configDownload** commands in the *Fabric OS Command Reference*.

Adding a New Switch or Fabric

When a new switch is added to the fabric, it automatically takes on the zone configuration information from the fabric. Use the **cfgActvShow** command to verify that the zoning information is the same on each switch in the fabric

If you are adding a switch that is already configured for zoning, use the **cfgClear** and **cfgSave** commands (or use **cfgClear** and **cfgDisable** if there is an effective configuration) before connecting it to the zoned fabric.

Adding a new fabric that has no zone configuration information to an existing fabric is very similar to adding a new switch. All switches in the new fabric inherit the zoning configuration data. If a zone configuration is in effect, then the same configuration becomes the enabled configuration. The **cfgActvShow** command will display the same information on all switches in the newly formed fabric.

Before the new fabric can merge successfully, it must pass the following criteria:

- Before merging zones

  To facilitate merging, check the following before merging switches or fabrics:

  - **Zoning licenses:** All switches must have a Zoning license enabled.
  - **Native operating mode**: All switches must be in the native operating mode.
  - **Secure Fabric OS**: If one switch has Brocade Secure Fabric OS enabled, all switches in the fabric must have Brocade Secure Fabric OS. See the *Secure Fabric OS Administrator's Guide* for more information.
  - **Default Zone:** The switch being merged into the existing fabric should be configured with the same default zone mode as the existing switches**.**

- Merging and segmentation

  The fabric is checked for segmentation during power-up or when a switch is disabled or enabled, or when a new switch is added.

  The database is the zone configuration database. This is the data displayed as the "defined configuration" in the **cfgShow** command. It is stored in nonvolatile memory by the **cfgSave** command. This database is a replicated database, which means that all switches in the fabric will have copy of this database. When a change is made to the defined configuration, the switch where the changes were made must close its transaction for the change to get propagated throughout the fabric.

  If you have implemented default zoning you must set the switch you are adding into the fabric the same default zone mode setting as the rest of the fabric to avoid segmentation.

- Merging rules

  Observe these rules when merging zones:

  - Local and adjacent configurations: If the local and adjacent zone database configurations are the same, they will remain unchanged after the merge.
  - Effective configurations: If there is an effective configuration between two switches, the zone configuration in effect must match.
  - Zone object naming: If a zoning object has the same name in both the local and adjacent defined configurations, the object types and member lists must match. When comparing member lists, the content and order of the members are important.
  - Objects in adjacent configurations: If a zoning object appears in an adjacent defined configuration, but not in the local defined configuration, the zoning object is added to the local defined configuration. The modified zone database must fit in the nonvolatile memory area allotted for the zone database.
  - Local configuration modification: If a local defined configuration is modified because of a merge, the new zone database is propagated to other the switches within the merge request.

- Merging Two Fabrics

  Both fabrics have identical zones and configurations enabled, including the default zone mode. The two fabrics will join to make one larger fabric with the same zone configuration across the newly created fabric.

If the two fabrics have different zoning configurations, they will be merged. If the two fabrics cannot join, the ISL between the switches will segment.

- Merge Conflicts

  When a merge conflict is present, a merge will not take place and the ISL will segment. Use the **switchShow** or **errLogShow** commands to obtain additional information about possible merge conflicts, because many non-zone related configuration parameters can cause conflicts

  If the fabrics have different zone configuration data, the system attempts to merge the two sets of zone configuration data. If the zones cannot merge, the ISL will be segmented.

  A merge is not possible if any of the following conditions exist:

  - Configuration mismatch: Zoning is enabled in both fabrics and the zone configurations that are enabled are different in each fabric.
  - Type mismatch: The name of a zone object in one fabric is used for a different type of zone object in the other fabric.
  - Content mismatch: The definition of a zone object in one fabric is different from the definition of zone object with the same name in the other fabric.
  - Zone Database Size: If the zone database size exceeds the maximum limit of another switch.

Note

If the zoneset members on two switches are not listed in the same order, the configuration is considered a mismatch, resulting in the switches being segmented from the fabric. For example: `cfg1 = z1; z2` is different from `cfg1 = z2; z1,` even though members of the configuration are the same. If zoneset members on two switches have the same names defined in the configuration, make sure zoneset members are listed in the same order.

Splitting a Fabric

If the connections between two fabrics are no longer available, the fabric will segment into two separate fabrics. Each new fabric will retain the same zone configuration.

If the connections between two fabrics are replaced and no changes have been made to the zone configuration in either of the two fabrics, then the two fabrics will merge back into one single fabric. If any changes that cause a conflict have been made to either zone configuration, then the fabrics might segment.

Using Zoning to Administer Security

Zones provide controlled access to fabric segments and establish barriers between operating environments. They isolate systems with different uses, protecting individual systems in a heterogeneous environment; for example, when zoning is in secure mode, no merge operations occur.

Brocade Advanced Zoning is configured on the primary Fabric Configuration Server (FCS). The primary FCS switch makes zoning changes and other security-related changes. The primary FCS switch also distributes zoning to all other switches in the secure fabric. All existing interfaces can be used to administer zoning (depending on the policies; see the *Secure Fabric OS Administrator's Guide* for information about security policies).

You must perform zone management operations from the primary FCS switch using a zone management interface, such as telnet or Advanced Web Tools. You can alter a zoning database, provided you are connected to the primary FCS switch.

When two secure fabrics join, the traditional zoning merge does not occur. Instead, a zoning database is downloaded from the primary FCS switch of the merged secure fabric. When E_Ports are active between two switches, the name of the FCS server and a zoning policy set version identifier are exchanged between the switches. If the views of the two secure fabrics are the same, the fabric's primary FCS server downloads the zoning database and security policy sets to each switch in the fabric. If there is a view conflict, the E_Ports are segmented due to incompatible security data.

As part of zoning architecture, you must determine which of the two basic zoning architectures (hard or soft) works best for your fabric. With time and planning, the basic hard zone configuration works for most sites. If a site has additional security needs, use the additional layer of Secure Fabric OS, apart from the standard zoning architecture.

Note
Secure Fabric OS requires the activation of a Brocade security license and an Advanced Zoning license.

Resolving Zone Conflicts

Zone conflicts can be resolved by saving a configuration file with the **configUpload** command, examining the zoning information in the file, and performing a cut and paste operation so that the configuration information matches in the fabrics being merged.

After examining the configuration file, you can choose to resolve zone conflicts by using the **cfgClear** command followed by the **cfgDisable** command on the incorrectly configured segmented fabric, followed by a **portDisable**/**portEnable** command on one of the ISL ports that connects the fabrics. This will cause a merge, making the fabric consistent with the correct configuration.

Caution
Be careful using the **cfgClear** command, because it deletes the defined configuration.

Table 20-14 lists considerations for zoning architecture.

Table 20-14 Considerations for Zoning Architecture

Item	Description
Type of zoning: hard or soft (session-based)	If security is a priority, hard zoning is recommended.
Use of aliases	The use of aliases is optional with zoning. Using aliases requires structure when defining zones. Aliases will aid administrators of zoned fabric in understanding the structure and context.
Security requirements	Evaluate the security requirements of the fabric. If additional security is required, add Brocade Secure Fabric OS into the fabric.
Interoperability Fabric	If the fabric includes a third-party switch product, only WWN zoning is supported. Other types of zoning, including QuickLoop, are not supported.
QLFA zones	Evaluate if the fabric will have QuickLoop Fabric Assist (QLFA) or QuickLoop (QL) in it, and consider the following items before creating and setting up QLFA zones: QuickLoop Zoning—QuickLoop/QuickLoop zones cannot run on Fabric OS v4.1.0 or later. However, Fabric OS can manage (create, remove, update) QL zones. QuickLoop Fabric Assist—A switch running Fabric OS v4.1.0 or later cannot have a Fabric Assist host directly connected to it. However, such a switch can be part of a Fabric Assist zone if a Fabric Assist host is connected to a compatible switch in the fabric.
Testing	Testing a (new) zone configuration. Before implementing a zone, you should run the Zone Analyzer from Web Tools to isolate any possible problems. This is especially useful as fabrics increase in size.
Effect of changes in a production fabric	Zone changes in a production fabric can result in a disruption of I/O under conditions where an RSCN is issued as a result of a zone change and the HBA is unable to process the RSCN fast enough. Though RSCNs are a normal part of a functioning SAN, the pause in I/O might not be acceptable. For these reasons, it is recommended to perform zone changes only when the resulting behavior is predictable and acceptable. Changing HBA drivers can shorten the response time in relation to the RSCN.
Confirming operation	After changing or enabling a zone configuration, you should confirm that the nodes and storage are able to identify and access one another. Depending on the platform, you might need to reboot one or more nodes in the fabric with the new changes.

Configuring and Monitoring FCIP Tunneling

Chapter 21

This chapter describes the concepts and procedures for connecting and monitoring the performance of SAN fabrics over IP-based networks using the optional Brocade Fibre Channel over IP (FCIP) Tunneling Service feature and the Brocade WAN analysis tool. The following topics are contained in this chapter:

- "About FCIP Tunneling" on page 21-1
- "FCIP Licensing" on page 21-2
- "Tunneling and IPSec" on page 21-6
- "Configuring FCIP Tunnels" on page 21-11
- "Fastwrite and Tape Pipelining" on page 21-16
- "Checklist for Configuring FCIP Links" on page 21-25
- "Troubleshooting FCIP Links" on page 21-25
- "WAN Performance Analysis Tools" on page 21-26

Supported Platforms for FCIP

Fabric OS supports FCIP ISLs between two SilkWorm switches (SilkWorm 7500 or SilkWorm 48000 with an FR4-18i blade).

The GbE ports on the FC4-16IP do not support FCIP tunneling.

Note

The Fabric OS v5.2.0 FCIP Tunneling Service is not compatible with the XPath FCIP service; nor is it compatible with any other vendor's implementation.

About FCIP Tunneling

The Fibre Channel over IP (FCIP) Tunneling Service is an optional feature that enables you to use Fibre Channel "tunnels" to connect SANs over IP-based networks. An FCIP tunnel uses TCP to transport data between a pair of Fibre Channel switches. Although there can be more than one TCP connection between the pair of Fibre Channel switches, from the Fibre Channel fabric point of view, the FCIP tunnel is just a connection between switches. This means that all IP network and protocols remain invisible.

FCIP tunneling enables you to connect fabrics over an IP connection. Depending on the port type you select, the fabrics can merge or remain separate while sharing devices. For example, you could connect one central office to different branch offices without having to merge the fabrics.

The port types for FCIP tunneling are either VE_Port or VEX_Port.

- An FCIP tunnel using VE_Ports will merge the two fabrics.
- An FCIP tunnel using a VEX_Port will not merge the fabrics. A VEX_Port can connect only to a VE_Port.

The IP encapsulation of the Fibre Channel frame on one port and the reconstruction of Fibre Channel frames on the other port is transparent to the initiator and target.

The tunnels on each GbE port are numbered 0 through 7, and each tunnel corresponds to a physical GbE port. For example:

- Tunnel 0 on GbE port 1 corresponds to logical FC port 24.
- Tunnel 0 on GbE port 0 corresponds to logical FC port 16.
- Tunnel 1 on GbE port 1 corresponds to logical FC port 25.

The GbE physical ports enable you to configure the following tunnels and corresponding virtual ports. Note that ports 16 through 31 are virtualized ports.

Table 21-1 Tunnels and Virtual Port Numbering Once Configured

Physical Port	Tunnels (portShow fcipTunnel slot/ge)	Virtual ports (switchShow)
ge0	0–7	16–23 (area numbers might be different depending on the slot chosen for the FR4-18i blade)
ge1	0–7	24–31 (area numbers might be different depending on the slot chosen for the FR4-18i blade)

FCIP also supports:

- Configuration and management of GbE ports and the virtual ports, IP interfaces, and tunnels enabled by GbE ports
- Compression and decompression of Fibre Channel frames moving through FCIP tunnels; by default, the compression mode of Fibre Channel frames moving through FCIP tunnels is *off*
- Traffic shaping that adheres to a rate limit on a per tunnel basis
- FCIP tunnel/GbE port event notification
- Fibre Channel routing service capabilities over FCIP tunnels

FCIP Licensing

To use FCIP, you must first install the Brocade FCIP license. Use the **licenseShow** command to verify the license is present on both switches used in the FCIP tunnel.

IPSec requires a license in addition to the FCIP license.

For details on obtaining and installing licensed features, see “Maintaining Licensed Features” on page 2-14.

FCIP Tunneling Concepts

FCIP tunneling uses the following terms and concepts (also see Figure 21-1):

- Tunnel

  An FCIP tunnel carries Fibre Channel traffic (frames) over IP networks such that the Fibre Channel fabric and all Fibre Channel devices in the fabric are unaware of the presence of the IP network. Fibre Channel frames “tunnel” through IP networks by fragmenting frames, encapsulating the result in IP packets on entering the tunnel, and then reassembling them as they leave the tunnel.

- GbE ports

  Two GbE ports are available on the SilkWorm 7500 and SilkWorm 48000 with an FR4-18i blade, for a total of four GbE ports. These ports support the FCIP feature with link speeds up to 1 Gbit/sec. Each GbE port (ge0, ge1) supports up to eight FCIP tunnels, for a total of sixteen virtual ports that can be configured as either VE_Ports, or VEX_Ports.

  Note

  The ports on the SilkWorm 7500 and FR4-18i are initially persistently disabled. See “Enabling Persistently Disabled Ports” on page 21-12 for information on enabling these ports. You should configure FCIP and the FC-FC Routing Service before enabling the ports. See Chapter 11, “Using the FC-FC Routing Service”.

- VE_Port

  Special types of ports, called VE_Ports (virtual E_Port), function somewhat like an E_Port. The link between an VE_Port and a VE_Port is called an “interswitch link” (ISL).

  - You can configure multiple ISLs from a SilkWorm 7500 or SilkWorm 48000 with an FR4-18i blade. After you configure the VE_Ports on either two SilkWorm 7500s or SilkWorm 48000s with the FR4-18i blade, an FCIP connection is established between them.
  - A VE_Port can be connected only to a VE_Port. VE_Port to VEX_Port tunnels are not supported.
  - VE_Ports do not prevent fabrics from merging.

- VEX_Port

  A VEX_Port enables FC-FC Routing Service functionality over an FCIP tunnel. VEX_Ports enable *interfabric links* (IFLs) so that you can share storage and servers between fabrics without merging those fabrics.

  - You can have up to eight VEX_Ports per GbE on the SilkWorm 7500 or SilkWorm 48000 with an FR4-18i blade.
  - A VEX_Port can only be connected to a VE_Port. VEX_Port to VEX_Port tunnels are not supported.
  - There is no trunking through GbE ports that have been configured as VE_Ports or VEX_Ports. See “EX_Port Frame Trunking (Optional)” on page 11-30 for trunking details.

Note

In Figure 21-1, because FCIP was configured with VE_Ports, the switches will merge over the IP WAN to become a single fabric. If any of the VE_Ports had been configured as VEX_Ports, that portion of the fabric would remain a separate fabric, but still enable sharing of storage and server devices.

Figure 21-1, illustrates a portion of a Fibre Channel network using FCIP. The FCIP interswitch link (VE_Ports connected over the IP WAN network) joins the office and data center SANs into a single larger SAN.

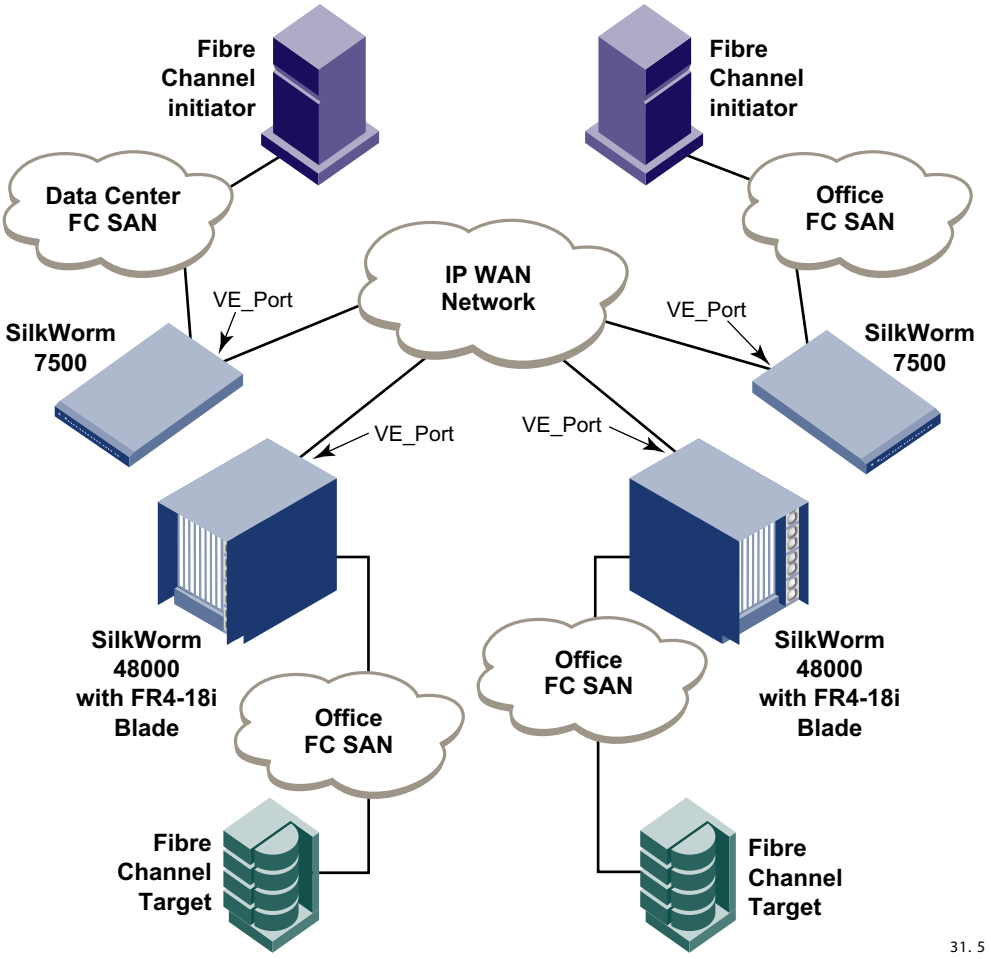

Figure 21-1 Network Using FCIP

Port Numbering

Port numbering differs on different hardware platforms.

Port Numbering on the FR4-18i Blade

There are 16 physical Fibre Channel ports and 2 physical GbE ports on the FR4-18i blade. The 2 GbE ports (ge0 and ge1) support up to 8 FCIP tunnels each (each FCIP tunnel is represented and managed as a VE_Port or VEX_Port).

- Ports 0–15 correspond to the physical Fibre Channel ports.
- Ports 16–23 are logical Fibre Channel ports on the physical GbE port, ge0.

- Ports 24–31 correspond to logical Fibre Channel ports on physical GbE port, ge1 (see Figure 21-2).

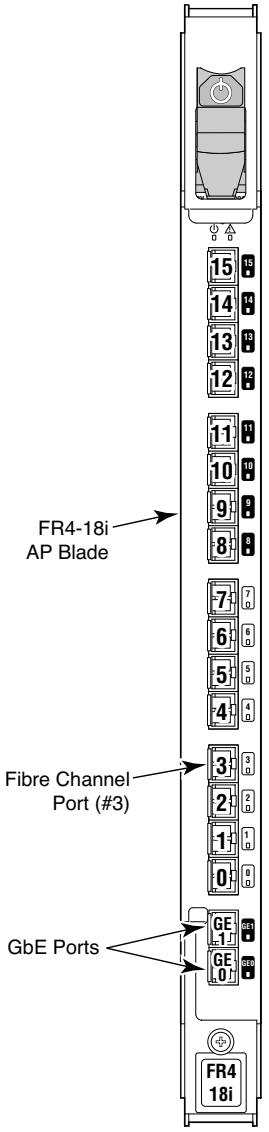

Figure 21-2 FR4-18i Port Numbering

You manage the FR4-18i blade as if it has 32 Fibre Channel ports (16 standard Fibre Channel ports, and 16 virtual Fibre Channel Ports). Specify port addresses using the slot and port numbers. For example, to disable VE_Port 18 on slot 1, use **portDisable 1/18**. To disable GbE port 1 on slot 1,use **portDisable 1/ge1**.

Port Numbering on the SilkWorm 7500

You do not need to specify slot numbers for the SilkWorm 7500. You see the GbE ports as ge0 and ge1, and the Fibre Channel ports are numbered 0–15. Moving from left to right on the front of the chassis, are the 16 Fibre Channel ports, followed by the 2 GbE ports.

You manage the SilkWorm 7500 as if it had 32 Fibre Channel ports (16 standard Fibre Channel ports, and 16 virtual Fibre Channel Ports) and 2 GbE ports. Specify port addresses using the slot and port numbers. For example, to disable VE_Port 18 on slot 1, use **portDisable 1/18**. To disable GbE port 1 on slot 1, use **portDisable 1/ge1**.

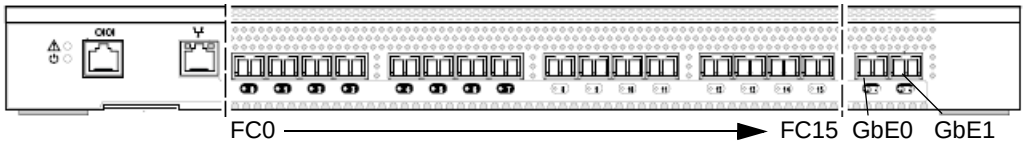

Figure 21-3 SilkWorm 7500 Port Numbering

Tunneling and IPSec

Internet Protocol security (IPSec) uses cryptographic security to ensure private, secure communications over Internet Protocol networks. IPSec supports network-level data integrity, data confidentiality, data origin authentication, and replay protection. It helps secure your SAN against network-based attacks from untrusted computers, attacks that can result in the denial-of-service of applications, services, or the network, data corruption, and data and user credential theft. By default, when creating an FCIP tunnel, IPSec is disabled.

FCIP tunneling with IPSec enabled will support maximum throughput as follows:

- Unidirectional—approximately 104MB/sec
- Bidirectional—approximately 90MB/sec

Used to provide greater security in tunneling on an FR4-18i blade or a Brocade SilkWorm 7500 switch, the IPSec feature does not require you to configure separate security for each application that uses TCP/IP. When configuring for IPSec, however, you must ensure that there is an FR4-18i blade or a Brocade SilkWorm 7500 switch in each end of the FCIP tunnel. IPSec works on FCIP tunnels with or without IP compression (IPComp).

IPSec requires an IPSec license in addition to the FCIP license.

IPSec uses some terms that you should be familiar with before beginning your configuration. These are standardized terms, but are included here for your convenience.

Table 21-2 IPSec Terminology

Term	Definition
AES	Advanced Encryption Standard. FIPS 197 endorses the Rijndael encryption algorithm as the approved AES for use by US Government organizations and others to protect sensitive information. It replaces DES as the encryption standard.
AES-XCBC	Cipher Block Chaining. A key-dependent one-way hash function (MAC) used with AES in conjunction with the Cipher-Block-Chaining mode of operation, suitable for securing messages of varying lengths, such as IP datagrams.
AH	Authentication Header - like ESP, AH provides data integrity, data source authentication, and protection against replay attacks but does not provide confidentiality.
DES	Data Encryption Standard is the older encryption algorithm that uses a 56-bit key to encrypt blocks of 64-bit plain text. Because of the relatively shorter key length, it is not a secured algorithm and no longer approved for Federal use.
3DES	Triple DES is a more secure variant of DES, it uses 3 different 56-bit keys to encrypt blocks of 64-bit plain text. The algorithm is FIPS-approved for use by Federal agencies.
ESP	Encapsulating Security Payload is the IPSec protocol that provides confidentiality, data integrity and data source authentication of IP packets, and protection against replay attacks.
MD5	Message Digest 5, like SHA-1, is a popular one-way hash function used for authentication and data integrity.
SHA	Secure Hash Algorithm, like MD5, is a popular one-way hash function used for authentication and data integrity.
MAC	Message Authentication Code is a key-dependent, one-way hash function used for generating and verifying authentication data.
HMAC	A stronger MAC because it is a keyed hash inside a keyed hash.
SA	Security association is the collection of security parameters and authenticated keys that are negotiated between IPSec peers.

The following limitations apply to using IPSec:

- IPv6, NAT, and AH are not supported.
- You can only create a single secure tunnel on a port; you cannot create a nonsecure tunnel on the same port as a secure tunnel.
- IPSec specific statistics are not supported.
- Fastwrite and tape pipelining (see page 21-16) cannot be used in conjunction with secure tunnels.
- To change the configuration of a secure tunnel, delete the tunnel and re-create it with the desired options.
- Jumbo frames are not supported for IPSec.
- There is no RAS message support for IPSec.
- Only a single route is supported on an interface with a secure tunnel.

Configuring IPSec

IPSEC requires predefined configurations for IKE and IPSEC. You can enable IPSEC only when these configurations are well-defined and properly created in advance.

The following steps provide an overview of the IPSec protocol. All of these steps require that the correct policies have been created. Because policy creation is an independent procedure from FCIP tunnel creation, you must know which IPSec configurations have been created. This ensures that you choose the correct configurations when you enable an IPSEC tunnel.

- Some traffic from an IPSec peer with the lower local IP address initiates the IKE negotiation process.
- IKE negotiates SAs and authenticates IPSec peers during phase 1 that sets up a secure channel for negotiation of phase 2 (IPSec) SAs.

  IKE negotiates SA parameters, setting up matching SAs in the peers. Some of the negotiated SA parameters include encryption and authentication algorithms, Diffie-Hellman group and SA lifetimes.
- Data is transferred between IPSec peers based on the IPSec parameters and keys stored in the SA database.
- IPSec tunnel termination. SA lifetimes terminate through deletion or by timing out.

The first step to configuring IPSec is to create a policy for IKE and a policy for IPSec. Once the policies have been created, you assign the policies when creating the FCIP tunnel.

IKE negotiates SA parameters and authenticates the peer using the preshared key authentication method. Once the 2 phases of the negotiation are completed successfully, the actual encrypted data transfer can begin.

IPSec policies are managed using the **policy** command.

You can configure up to 32 IKE and 32 IPSec policies. Policies cannot be modified; they must be deleted and recreated in order to change the parameters. You can delete and recreate any policy as long as the policy is not being used by an active FCIP tunnel.

Each FCIP tunnel is configured separately and may have the same or different IKE and IPSec policies as any other tunnel. Only one IPSec tunnel can be configured for each GbE port.

IPSec Parameters

When creating policies, the following parameters are fixed and cannot be modified:

Table 21-3 Fixed policy parameters

Parameter	Fixed Value
IKE negotiation protocol	Main mode
ESP	Tunnel mode
IKE negotiation authentication method	Preshared key
3DES encryption	Key length of 168 bits
AES encryption	Key length of 128 or 256

The following parameters can be modified:

Table 21-4

Parameter	Description
Encryption Algorithm	3DES—168-bit key AES-128—128-bit key (default) AES-256—256-bit key
Authentication Algorithm	SHA-1—Secure Hash Algorithm (default) MD5—Message Digest 5 AES-XCBC—Used only for IPSec
Security Association lifetime in seconds	The lifetime in seconds of the security association. If PFS is enabled, a new IKE SA using new key material will be negotiated before this value expires. Default is 28800 sec.
PFS (Perfect Forward Secrecy)	Applies only to IKE policies. Choices are On/Off and default is On.
Diffie-Hellman group	Group 1—768 bits (default) Group 14—2048 bits

Managing Policies

Use the **policy** command to create, delete, and show IKE and IPSec policies.

To create a new policy

1. Log in to the switch as admin.
2. At the command prompt, type:

```
policy --create type number [-enc encryption_method][-auth
authentication_algorithm] [-pfs off|on] [-dh DH_group] [-seclife secs]
```

where:

type and *number* — The type of policy being created (IKE or IPSec) and the number for this type of policy. To easily determine how many policies have been created, consider using sequential numbering. The range of valid values is any whole number from 1 through 32.

encryption_method — The supported type of encryption. Valid options are 3DES, AES-128, and AES-256. AES-128 is the default.

authentication_algorithm — The authentication algorithm. Valid options are SHA-1, MD5, and AES-XCBC (IPSec only). HA-1 is the default.

DH_Group — The Diffie-Hellman group. Supported groups are Group 1 and Group 14. Group 1 is the default.

secs — The security association lifetime in seconds. 28800 is the default.

The following example shows how to create IKE policy number 10 using 3DES encryption, MD5 authentication, and Diffie-Hellman Goup 1:

```
switch:admin06> policy --create ike 10 -enc 3des -auth md5 -dh 1
The following policy has been set:

IKE Policy 10
----------------------------------------
Authentication Algorithm: MD5
Encryption: 3DES
Perfect Forward Secrecy: on
Diffie-Hellman Group: 1
SA Life (seconds): 28800

Operation Succeeded
```

For a complete description of the **policy** command, see the*Fabric OS Command Reference*.

To display policy settings

1. To display the settings for a single policy, type:

   `policy --show` *`type number`*

   For example, to view the IPSec 1 policy, type:

   `policy --show ipsec 1`

2. To display the policy settings for all defined policies, type:

   `policy --show` *`type`* `all`

   The example below shows all of the IKE policies defined; in this example, there are two IKE policies.

```
switch:admin06> policy --show ike all
IKE Policy 1
----------------------------------------
Authentication Algorithm: MD5
Encryption: UNKNOWN
Perfect Forward Secrecy: off
Diffie-Hellman Group: 1
SA Life (seconds): 0


IKE Policy 32
----------------------------------------
Authentication Algorithm: SHA-1
Encryption: AES-128
Perfect Forward Secrecy: on
Diffie-Hellman Group: 1
SA Life (seconds): 28800

Operation Succeeded
```

Policies cannot be modified. You must delete and then recreate a policy with the newly-determined parameters.

To delete a policy

1. Log in to the switch as admin.
2. At a command prompt, type:

   policy --delete *type number*

   where *type* is the policy type and *number* is the number assigned.

   For example, to delete the IPSec policy number 10:

```
switch:admin06> policy --delete ipsec 10
The policy has been successfully deleted.
```

Configuring FCIP Tunnels

You can create only one FCIP tunnel on a given pair of IP address interfaces (local and remote). You can create multiple FCIP tunnels on a single IP interface if either the local or remote IP interface is unique and does not have any other FCIP tunnel on it. When the GbE port has a valid SFP and is physically connected to any other GbE port, the status output from the **switchShow** command is online.

Due to an IPSec RASlog limitation, you may not be able to determine an incorrect configuration that causes an IPSec tunnel not to become active. This misconfiguration can occur on either end of the tunnel. As a result, you must correctly match the encryption method, authentication algorithm, and other configurations for on each end of the tunnel.

Note
The procedures in this section demonstrate configuring FCIP tunnels for remote switches. The same procedures apply to local switches and need to be performed there as well.

See the *Fabric OS Command Reference* for detailed information on using the commands in this section.

Following are the steps for configuring an FCIP tunnel:

1. “Enabling Persistently Disabled Ports” on page 21-12
2. “Defining the IP Interface of Each Virtual Port” on page 21-13
3. “Configuring the GbE Ports” on page 21-13
4. “Adding IP Routes on a GbE Port” on page 21-14
5. “Verifying IP Connectivity” on page 21-15
6. “Verifying the FCIP Tunnel Configuration” on page 21-22

Before you begin configuring FCIP tunnels, verify that you have an FCIP license installed. See “Maintaining Licensed Features” on page 2-14.

Enabling Persistently Disabled Ports

Ports on the SilkWorm 7500 and FR4-18i are disabled by default. Before you can configure FCIP tunnels, you must persistently enable the ports.

Caution
VEX_Port Users: If the fabric is already connected, you must leave the ge0 and ge1 ports disabled until *after you have configured the VEX_Port;* this will prevent unintentional merging of the two fabrics.

To enable a persistently disabled port

1. Enter the **portCfgShow** command to view ports that are persistently disabled.
2. After identifying the ports, enter the **portCfgPersistentEnable** command to enable the ports.
3. Disable the port during FCIP configuration by entering the **portDisable** [*slot*/]*port* command.
4. Enter the **portCfgShow** command to verify the port is persistently enabled as shown below:

```
switch:admin06> portcfgpersistentenable 8/16
switch:admin06> portcfgpersistentenable 8/17
switch:admin06> portcfgpersistentenable 8/18
switch:admin06> portcfgpersistentenable 8/19

switch:admin06> portcfgshow
Ports of Slot 8    0  1  2  3    4  5  6  7    8  9 10 11   12 13 14 15
-----------------+--+--+--+--+----+--+--+--+----+--+--+--+----+--+--+--
Speed             AN AN AN AN   AN AN AN AN   AN AN AN AN   AN AN AN AN
Trunk Port        ON ON ON ON   ON ON ON ON   ON ON ON ON   ON ON ON ON
Long Distance     .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
VC Link Init      .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Locked L_Port     .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Locked G_Port     .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Disabled E_Port   .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
ISL R_RDY Mode    .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
RSCN Suppressed   .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Persistent DisableON ON ON ON   ON ON ON ON   ON ON ON ON   ON ON ON ON
NPIV capability   ON ON ON ON   ON ON ON ON   ON ON ON ON   ON ON ON ON
EX Port           .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Mirror Port       .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..

Ports of Slot 8   16 17 18 19   20 21 22 23   24 25 26 27   28 29 30 31
-----------------+--+--+--+--+----+--+--+--+----+--+--+--+----+--+--+--
Speed             AN AN AN AN   AN AN AN AN   AN AN AN AN   AN AN AN AN
Trunk Port        ON ON ON ON   ON ON ON ON   ON ON ON ON   ON ON ON ON
Long Distance     .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
VC Link Init      .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Locked L_Port     .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Locked G_Port     .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Disabled E_Port   .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
ISL R_RDY Mode    .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
RSCN Suppressed   .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Persistent Disable.. .. .. ..   ON ON ON ON   ON ON ON ON   ON ON ON ON
NPIV capability   ON ON ON ON   ON ON ON ON   ON ON ON ON   ON ON ON ON
EX Port           .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..

                  where AN:AutoNegotiate, ..:OFF, ??:INVALID.
                        LM:L0.5
switch:admin06>
```

Defining the IP Interface of Each Virtual Port

Specify the IP interface on the virtual port using the **portCfg** command. You can define up to eight IP interfaces per GbE port. The syntax for this command is as follows:

portcfg ipif [*slot*/][**ge**]*port args*:

To create IP interfaces:

create *ipaddr netmask mtu_size*

To delete IP interfaces:

delete *ipaddr*

Use the **portShow** command to display all IP interfaces on a particular GbE port. Note that fabrics connected through FCIP are merged (unless they are connected as VEX_Ports).

To display IP interfaces on the specified GbE port:

portshow ipif [*slot*/][**ge**]*port*

The following example shows IP interfaces defined for slot 8 on GbE port ge0:

```
switch:admin06> portshow ipif 8/ge0
Port: 8/ge0
Interface       IP Address      NetMask         MTU
-----------------------------------------------------------
  0             192.168.100.40  255.255.255.0   1500
  1             192.168.100.41  255.255.255.0   1500
```

```
switch:admin06> portcfg ipif 8/ge0 create 192.168.100.40 255.255.255.0 1500
switch:admin06> portcfg ipif 8/ge0 create 192.168.100.41 255.255.255.0 1500
```

Then verify the newly-created IP interface using the **portShow** command:

```
switch:admin06> portshow ipif 8/ge0

Port: 8/ge0
Interface       IP Address      NetMask         MTU
-----------------------------------------------------------
  0             192.168.100.40  255.255.255.0   1500
  1             192.168.100.41  255.255.255.0   1500
```

Configuring the GbE Ports

Each GbE port has eight virtual channels. In order to enable FCIP on these ports you must configure them as either a VE_Port or a VEX_Port, depending on whether you want to the fabrics connected using FCIP to merge. Two fabrics connected with VE_Ports will merge into a single fabric. Two fabrics connected with a VEX_Port / VE_Port will not merge. The destination IFL port type for a VEX_Port should always be a VE_Port. By default the virtual ports will automatically become VE_Ports.

Use **portCfgVexport** to specify a port as a VEX_Port is as follows:

portcfgvexport [*slotnumber*/]*portnumber* [**-a** *admin*][**-f** *fabricid*][**-d** *domainid*] [**-p** *pidformat*][**-t** *fabric_parameter*]

where:

-a *admin*	Specify 1 to enable or 2 to disable the admin.
-f *fabricid*	Specify 1 to 128 for the fabric ID.
-d *domainid*	Specify 1 to 239 for the preferred domain ID.
-p *pidformat*	Specify 1 for core, 2 for extended edge, and 3 for native port ID format.
-t *fabric_parameter*	Specify 1 to enable or 2 to disable negotiate fabric parameters.

For example, to configure a port as a VEX_Port for slot number 8 in port number 18, enable the admin, with fabric ID 2 and preferred domain ID 220:

```
switch:admin06> portcfgvexport 8/18 -a 1 -f 2 -d 220
```

Adding IP Routes on a GbE Port

After defining the IP interface of the remote switch, you can define destination routes on an interface. You can specify a maximum of 32 routes per GbE port.

The syntax to create IP routes is:

portcfg iproute [*slot*/][**ge**]*port args* **create** *ipaddr netmask gateway_router metric*

The following example shows two routes being added to an interface:

```
switch:admin06> portcfg iproute 8/ge0 create 192.168.11.0 255.255.255.0 192.168.100.1
1
switch:admin06> portcfg iproute 8/ge0 create 192.168.12.0 255.255.255.0 192.168.100.1
1
```

The syntax to delete IP routes is:

portcfg iproute [*slot*/][**ge**]*port args* **delete** *ipaddr netmask*

The gateway address must be on the same IP subnet as one of the port IP addresses.

The following example verifies that the two routes have been successfully created:

```
switch:admin06> portshow iproute 8/ge0

Slot: 8 Port: ge0
IP Address      Mask            Gateway         Metric  Flags
-----------------------------------------------------------------
192.168.100.0   255.255.255.0   192.168.100.40   0       Interface
192.168.100.0   255.255.255.0   192.168.100.41   0       Interface
192.168.11.0    255.255.255.0   192.168.100.1    1
192.168.12.0    255.255.255.0   192.168.100.1    1
```

Verifying IP Connectivity

After you add the IP addresses of the routes, enter the **portCmd ping** command to ping a destination IP address from one of the source IP interfaces on the GbE port and verify the Ethernet IP to IP connectivity. This verification also ensures that data packets can be sent to the remote interface. You can ping a connection only if both ports have IP interfaces set.

Use the **portCmd --ping** command to ping a destination IP address from one of the source IP interfaces on the GbE port. Following is the syntax for the command:

portcmd --ping [*slot-number/*] *port-number* **-s** <source -ip> **-d** <dest-ip>
[**-n** <num-requests>][**-q** <*type-of-service*>][**-t** <*ttl*>][**-w**<*wait-time*]<**-z**<*size*>]

where:

-s <*source_ip*> Specifies the source IP interface that originates the ping request.

-d <*destination_ip*> Specifies the destination IP address to which to target the ping request.

-n <num-requests> (Optional) Generates specified number of ping requests. The default is 4.

-q <*type-of-service*> (Optional) Sets the “tos” in the ping requests to the type of service. The default is 0 (zero). The value must be an integer in the range from 0 through 255.

-t <*ttl*> (Optional) Specifies the time to live. The default value is 100.

-w <*wait-time*> (Optional) Specifies the time to wait for the response of each ping request. This parameter is specified in milliseconds and the default value is 5000 milliseconds (5 sec). The maximum allowed wait time for ping is 9000 milliseconds (9 sec).

-z <*size*> (Optional) Specifies the size in bytes of the ping packet to use. The default size is 64 bytes. The total size including ICMP/IP headers (28 bytes without IP options) cannot be greater than IP MTU configured on the interface.

Following is an example of pinging an address from a particular IP interface:

```
switch:admin06> portcmd --ping ge0 -s 192.175.5.100 -d 192.175.5.200
Pinging 192.175.5.200 from ip interface 192.175.5.100 on 0/ge0 with 64 bytes of data
Reply from 192.175.5.200: bytes=64 rtt=1ms ttl=64
Reply from 192.175.5.200: bytes=64 rtt=0ms ttl=64
Reply from 192.175.5.200: bytes=64 rtt=0ms ttl=64
Reply from 192.175.5.200: bytes=64 rtt=1ms ttl=64

Ping Statistics for 192.175.5.200:
        Packets: Sent = 4, Received = 4, Loss = 0 (0 percent loss)
        Min  RTT = 0ms, Max RTT = 1ms Average = 0ms
```

See the *Fabric OS Command Reference* for details about using the **portCmd ping** command.

Testing End-to-End IP Path Performance

After installing the IP configuration on the FCIP port (e.g. IP address, route entries), you can test the end-to-end path performance characteristics of the IP interface using WAN analysis tools. For more information, see “WAN Performance Analysis Tools” on page 21-26. These tools enable you to monitor bandwidth, loss, round trip time, and other statistics for a traffic stream between a pair of Brocade port endpoints.

Fastwrite and Tape Pipelining

When the FCIP link is the slowest part of the network and it affects speed, consider using fastwrite and tape write acceleration, called “tape pipelining.” Supported only in Fabric OS 5.2.x and higher, fastwrite and tape pipelining are two features that provide accelerated speeds to FCIP tunnels in some configurations:

- Fastwrite accelerates the SCSI write I/Os over FCIP.
- Tape pipelining accelerates SCSI write I/Os to sequential devices (such as tape drives) over FCIP, which reduces the number of round-trip times needed to complete the I/O over the IP network and speeds up the process. In order to use tape pipelining, you must enable fastwrite as well.
- Both sides of an FCIP tunnel must have matching configurations for these features to work.

Unlike IPSec, compression, fastwrite, and tape pipelining features do not require any predefined configurations. This makes it possible to enable these features by adding optional parameters such as **–c**, **-f**, or **-t** when you create FCIP tunnels.

Enabling Fastwrite and Tape Pipelining

Fastwrite and tape pipelining require no parameters: both features are enabled by turning them on during the tunnel creation process. They are enabled on a per-FCIP tunnel basis. See“Configuring FCIP Tunnels” on page 21-11 for details.

Constraints for Fastwrite and Tape Pipelining

Consider the constraints described in Table 21-5 when configuring tunnels to use either of these features.

Table 21-5 Using Fastwrite and Tape Pipelining

Fastwrite	Tape Pipelining
Each GbE port supports up to 2048 simultaneous accelerated exchanges, which means *a total of 2048 simultaneous exchanges combined* for fastwrite and tape pipelining.	Each GbE port supports up to 2048 simultaneous accelerated exchanges, which means *a total of 2048 simultaneous exchanges combined* for fastwrite and tape pipelining.
Does not affect FICON traffic	Does not affect FICON traffic

Table 21-5 Using Fastwrite and Tape Pipelining (Continued)

Fastwrite	Tape Pipelining
Does not support multiple equal-cost path configurations (see "Supported Configurations").	Does not support multiple equal-cost path configurations or multiple non-equal-cost path configurations (see "Supported Configurations").
Class 3 traffic is accelerated with fastwrite.	Class 3 traffic is accelerated between host and sequential device.
	With sequential devices (tape drives), there are 1024 initiator-tape (IT) pairs per GbE port, but 2048 initiator-tape-LUN (ITL) pairs per GbE port. The ITL pairs are shared among the IT pairs. For example: • 2 ITL pairs for each IT pair as long as the target has two LUNs. • If a target has 32 LUNs, 32 ITL pairs for IT pairs. In this case, only 64 IT pairs are associated with ITL pairs. The rest of the IT pairs are not associated to any ITL pairs, so no tape pipelining is performed for those pairs. By default, only fastwrite-based acceleration is performed on the unassociated pairs.
	Does not support multiple non-equal-cost path between host and sequential device

Supported Configurations

To help understand the supported configurations, consider the configurations shown in the two figures below. In both cases, there are no multiple equal-cost paths. In the first figure, there is a single tunnel with fastwrite and tape pipelining enabled. In the second figure, there are multiple tunnels, but none of them create a multiple equal-cost path.

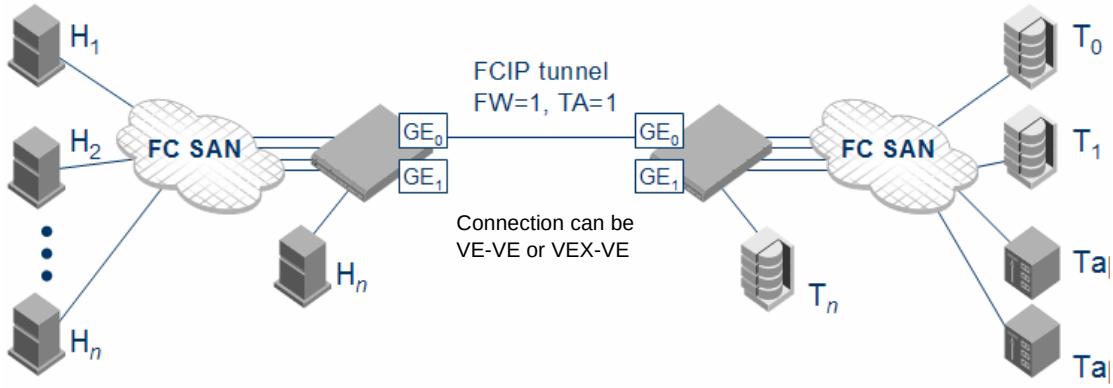

Figure 21-4 Single Tunnel, Fastwrite and Tape Pipelining Enabled

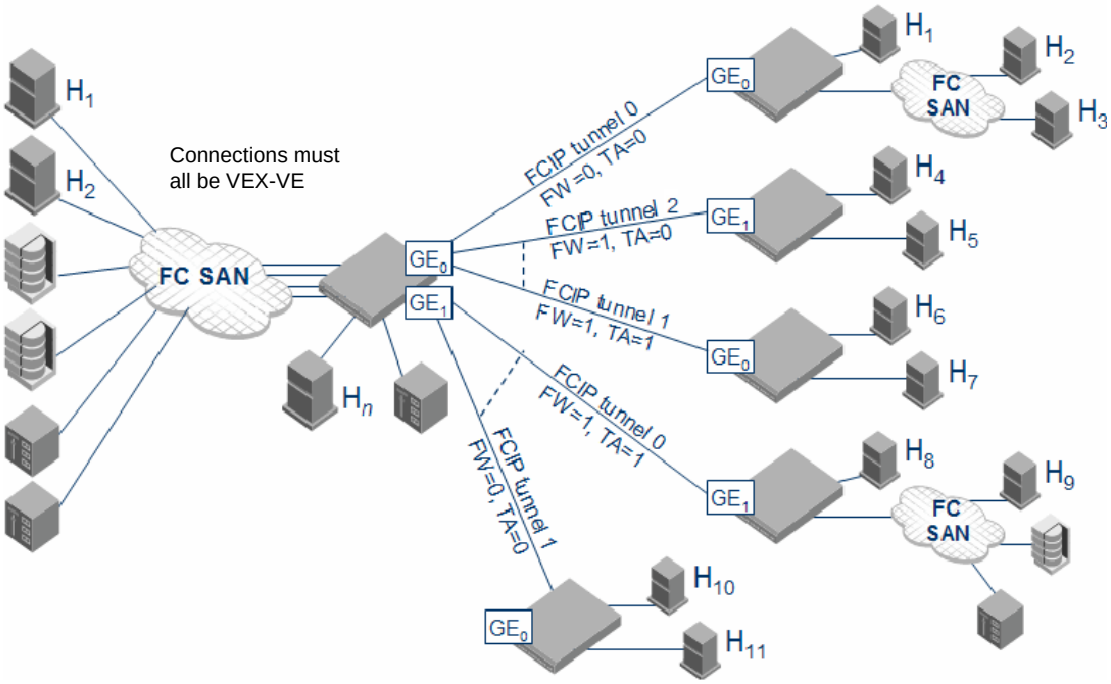

Figure 21-5 Multiple Tunnels to Multiple Ports, Fastwrite and Tape Pipelining Enabled on a Per-tunnel/Per-port Basis

Unsupported Configurations

The following configurations are not supported with fastwrite and tape pipelining. These configurations use multiple equal-cost paths.

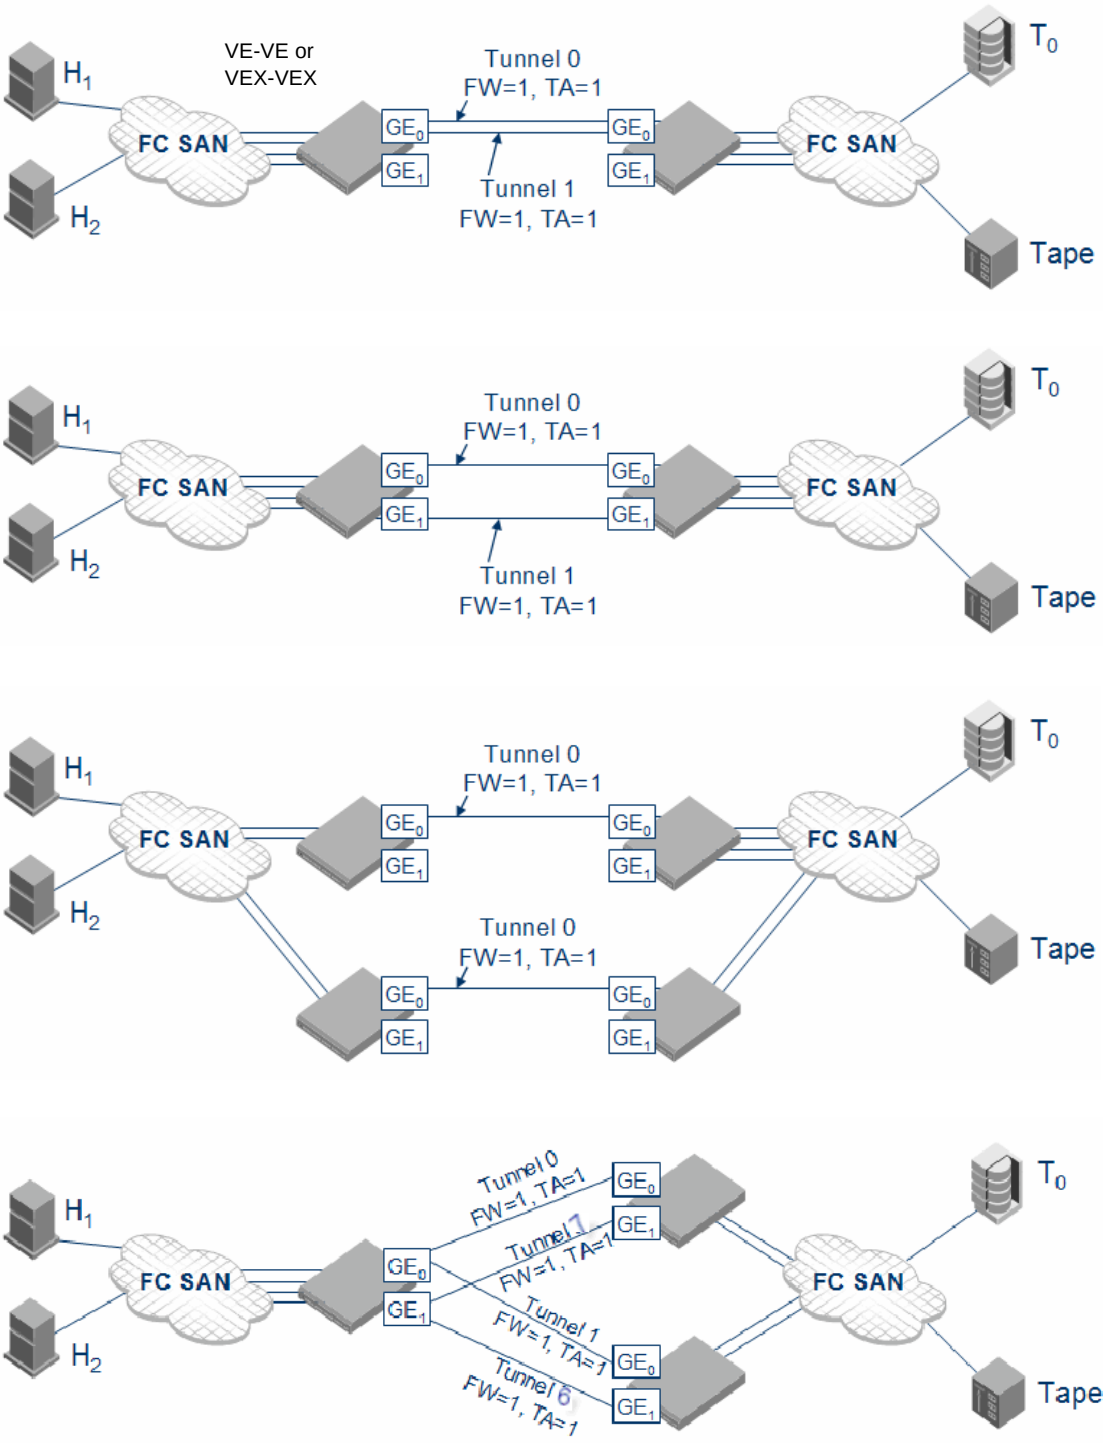

Figure 21-6 Unsupported Configurations with Fastwrite and Tape Pipelining

Configuring FCIP Tunnels

After you have verified licensing and connectivity between source and destination IP interfaces, you can configure FCIP tunnels. As you plan the tunnel configurations, be aware that uncommitted rate tunnels use a minimum of 1000 Kb/sec, up to a maximum of available uncommitted bandwidth on the GbE port. The total bandwidth available on a GbE port is 1 Gbit/sec. You can configure tunnels as bidirectional entities with different commit rates in both directions.

Note

You cannot create FCIP tunnels that connect a SilkWorm 7500 or SilkWorm 48000 with an FR4-18i blade to a SilkWorm Multiprotocol Router Model AP7420.

The syntax for configuring FCIP Tunnels is provided in the following sections. For more information about Fabric OS commands, see the *Fabric OS Command Reference*.

FCIP Tunnel Create Option

Following is the syntax for the **portCfg fcipTunnel** command:

```
portcfg fciptunnel [slot/][ge] port args [optional_args] create <tunnel_id>
<remote_ip_addr> <local_ip_addr> <comm_rate> [<-csf>] [-n remote_wwn] [<-k timeout>]
[<-r retransmissions>] [<-m time>] [<-t>] [-ike <ike_number>] [-ipsec <ipsec_number>]
[-key <preshared_key>]
```

where:

-c	Enables compression on this tunnel.
-s	Disables selective acknowledgement code (SACK) on the specified tunnel.
-f	Enables fastwrite.
-n *remote_wwn*	Specifies the remote-side FC entity WWN.
-k *timeout*	Specifies the keep-alive timeout in seconds. The range of valid values is 8 through 7,200 sec and the default is 10. If tape pipelining is enabled both the default and minimum values are 80 sec.
-r *retransmissions*	Specifies the maximum number of retransmissions on the existing FCIP tunnel. The range of valid values is 1 through 16. If tape pipelining is enabled, the number of retransmissions is calculated based on the minimum retransmit time to ensure that tunnel does not timeout before the host times out (about 80 sec). Note that the value specified must be greater than the calculated value.
-m *time*	Specifies the minimum retransmit time, in milliseconds. The range of valid values is 20 through 5,000 ms and the default is 100 ms.
-t	Enables tape pipelining. If tape pipelining is enabled, fastwrite must also be enabled.
-ike *ike_number*	Specifies the IKE policy number to be used for this FCIP tunnel.
-ipsec *ipsec_number*	Specifies the IPSec policy number to be used for this FCIP tunnel.
-key *preshared_key*	The preshared key to be used during IKE authentication up to a maximum length of 32 bytes. It can be a double quoted string of alphanumeric characters. The range of valid values is 12 through 32 bytes.

FCIP Tunnel Modify and Delete Options

Note

Using the tunnel **Modify** option disrupts traffic on the specified FCIP tunnel for a brief period of time.

Following is the syntax for the **portCfg fcipTunnel** command to modify FCIP tunnels (you must specify at least one characteristic to modify):

```
portcfg fciptunnel [slot/][ge]port args [optional_args] modify <tunnel_num>
[<-b comm._rate>] [<-c 0|1>] [<-f 0|1>] [<-k timeout>] [<-m time>]
[<-r retransmissions>] [<-s 0|1>]
```

where:

-b *comm._rate*	New committed traffic rate in Kb/sec on the existing FCIP tunnel.
-c 0|1	Disables (0)/Enables (1) compression on the existing FCIP tunnel.
-f 0|1	Enables (1)/Disables (0) fastwrite.
-k *timeout*	New keep-alive timeout on the existing FCIP tunnel. A range of valid values is 8 through 7200 sec. If tape pipelining is enabled, the default and minimum value is 80 sec.
-m *time*	New minimum retransmit time for the existing FCIP tunnel. A range of valid values is 20 through 5000 ms (milliseconds).
-r *retransmissions*	Maximum number of retransmissions on the existing FCIP tunnel. A range of valid values is 1 through 16. If tape pipelining is enabled, the number of retransmissions is calculated based on the minimum retransmit time to ensure that the tunnel does not time out before the host times out (about 80 sec). If you change this value, the value specified must be greater than the calculated value.
-s 0|1	SACK ON (1)/SACK OFF (0) on the existing FCIP tunnel.
-t 0|1	Enables (1)/Disables (0) tape pipelining on the existing tunnel. If tape pipelining is enabled, you must also enable fastwrite.

The following example shows two tunnels created on slot 8, port ge0; the first with an uncommitted bandwidth (0):

```
switch:admin06> portcfg fciptunnel 8/ge0 create 2 192.168.100.50 192.168.100.40 0
switch:admin06> portcfg fciptunnel 8/ge0 create 3 192.168.100.51 192.168.100.41
10000
```

The following example shows an FCIP tunnel created between local IP interface <192.168.131.124>:

```
switch:admin06> portcfg  fciptunnel 3/ge0 create 6 10.1.1.44   192.168.131.124
155000
```

Following is the syntax for the **portCfg fcipTunnel** command to delete FICP tunnels:

```
portcfg fciptunnel [slot/][ge]port args [optional_args] delete <tunnel_num>
```

The following example shows two tunnels deleted on slot 8, port ge0:

```
switch:admin06> portcfg fciptunnel 8/ge0 delete 6
switch:admin06> portcfg fciptunnel 8/ge0 delete 7
```

Verifying the FCIP Tunnel Configuration

After you have created local and remote FCIP configurations, use the **portEnable** [*slot*/]*port* command to enable the port.

It is recommended that you verify that the tunnel configuration operation succeeded using the **portShow fcipTunnel** command (be sure to specify the slot/port numbers and number of tunnels). Look at the "Status" field to verify that the tunnel is now "Active."

You can also use **portShow fcipTunnel** to display performance characteristics, when running WAN tools on a pair of Brocade port endpoints. For more information about using **portShow** for performance analysis, see "FCIP Tunnel Performance Characteristics" on page 21-30.

Following is the syntax for the **portShow fcipTunnel** command to display FCIP tunnels on a GbE port:

portshow fciptunnel [*slot*/][**ge**]*port* [*args*]

where:

all Displays all FCIP tunnels.

tunnel_id Displays the specified FCIP tunnel.

The following example shows an active tunnel with fastwrite enabled but tape pipelining off:

```
switch:admin06> portshow fciptunnel ge0 all
----------------------------------------
Tunnel ID 0
Remote IP Addr 10.0.10.224
Local IP Addr 10.0.10.225
Remote WWN Not Configured
Local WWN 10:00:00:05:1e:37:91:dd
Compression on
Fastwrite on
Tape Pipelining off
Uncommitted bandwidth, minimum of 1000 Kbps (0.001000 Gbps)
SACK on
Min Retransmit Time 100
Keepalive Timeout 10
Max Retransmissions 8
Status : Active
Uptime 21 seconds
```

If IPSec has been enabled and a policy added to the configuration, you will see the policy information under the status section of the output, as shown below. The policy information is visible only when IPSec is configured and is displayed with the information shown in the example above when the **portShow** command is issued.

```
IKE Policy 1
Authentication Algorithm: SHA-1
Encryption: 3DES
Perfect Forward Security: off
Diffie-Hellman Group: 1
SA Life (seconds): 3600
SA Life (MB): 4500
IPSec Policy 1
Authentication Algorithm: SHA-1
Encryption: 3DES
Perfect Forward Security: off
Diffie-Hellman Group: 1
SA Life (seconds): 3600
SA Life (MB): 4500
Pre-Shared Key 0x12345000000000000000000000000000000
```

After FCIP tunnels are created, the configuration is saved in a persistent database. At this point, all configured FCIP tunnels now appear in the fabric as VE_Ports.

To verify that a VE_Port or VEX_Port is online

1. Use the **switchShow** command to view and verify that the FCIP tunnel is online.

```
switch:admin06> portenable 8/18
switch:admin06> portenable 8/19
switch:admin06> switchshow
switchName:switch
switchType:42.2
switchState:Online
switchMode:Native
switchRole:Subordinate
switchDomain:4
switchId:fffc04
switchWwn:10:00:00:60:69:80:0d:bc
zoning:ON (LSAN001)
switchBeacon:OFF
blade3 Beacon:  OFF
blade4 Beacon:  OFF
blade8 Beacon:  OFF
FC Router:ON
FC Router BB Fabric ID:1

Area Slot Port Media Speed State
====================================
 32    3    0   id    N4   Online    F-Port  50:03:0d:30:0d:13:00:09
 33    3    1   id    N4   Online    F-Port  50:03:0d:30:0d:13:00:11
 34    3    2   id    N4   Online    F-Port  50:03:0d:30:0d:13:00:13
 35    3    3   id    N4   Online    F-Port  50:03:0d:30:0d:13:00:15
 36    3    4   id    N2   Online    F-Port  21:00:00:e0:8b:08:bd:20
<output truncated>
210    8   18 --    --   Online    VE-Port  50:00:51:e3:51:55:3f:1e "fcr_xd_3_16"
(downstream)
211    8   19 --    --   Online    VE-Port  50:00:51:e3:70:42:5f:76 "fcr_xd_5_17"
(downstream)
<output truncated>
223    8   31   --    --   Offline
       8  ge0   id    1G   Online
       8  ge1   id    1G   Online
```

Checklist for Configuring FCIP Links

Table 21-6 can be used as a checklist for creating FCIP links.

Table 21-6 Steps for Configuring FCIP Links

Step	Command
1. Enable persistently disabled ports.	**portcfgpersistentenable** [*slot*/]*port*
2. Disable the ports while performing the configuration.	**portdisable** [*slot*/]*port*
3. Configure the port type as VE_Port or VEX_Port for both ports for a tunnel.	**portcfgvexport** [*slot*/][**ge**] *port*
4. Configure the IP interface for both ports of a tunnel.	**portcfg ipif** [*slot*/][**ge**]*port* **create** *ipaddr netmask mtu_size*
5. Verify the IP interface for both ports of a tunnel.	**portshow ipif** [*slot*/][**ge**]*port*
6. Configure FCIP tunnels.	**portcfg fciptunnel** [*slot*/][**ge**]*port* **create** *tunnel_id remote_ipaddr local_ipaddr comm_rate*
7. Verify FCIP tunnels.	**portshow fciptunnel** [*slot*/][**ge**]*port* [**all**|*tunnel_num*]
8. Enable the ports.	**portenable** [*slot*/]*port*

Note

See the *Fabric OS Command Reference* for detailed information on these commands.

Troubleshooting FCIP Links

The following list contains information for troubleshooting FCIP links:

- When deleting FCIP links, you must delete them in the exact reverse order they were created. That is, delete first the tunnels, then the IP interfaces, and finally the port configuration. The IP route information is removed automatically at this point.
- IP addresses are retained by slot in the system. If FR4-18i blades are moved to different slots without first deleting configurations, errors can be seen when trying to reuse these IP addresses.
- The **portCmd ping** command only verifies physical connectivity. This command does not verify that you have configured the ports correctly for FCIP tunnels.
- One port can be included in multiple tunnels, but each tunnel must have at least one port that is unique to that tunnel.
- Ports at both ends of the tunnel must be configured correctly for a FCIP tunnel to work correctly. These ports can be either VE_Ports or VEX_Ports. A VEX_Port must be connected to a VE_Port.
- When configuring routing over an FCIP link for a fabric, the edge fabric will use VE_Ports and the backbone fabric will use VEX_Ports for a single tunnel.

- If an FCIP tunnel fails with the “Disabled (Fabric ID Oversubscribed)” message, the solution is to reconfigure the VEX_Port to the same Fabric ID as all of the other ports connecting to the edge fabric.

WAN Performance Analysis Tools

Introduced in Fabric OS v5.2.0, WAN analysis tools are designed to estimate the end-to-end IP path performance characteristics between a pair of Brocade FCIP port endpoints. WAN tools include the following commands and options:

- **portCmd ipPerf**—Characterizes end-to-end IP path performance between a pair of Brocade FCIP ports. You can use the WAN tool **ipPerf** only on the FR4-18i or SilkWorm 7500 FCIP ports running Fabric OS v5.2.0 or later software.
- **portShow fcipTunnel**—Displays performance statistics generated from the WAN analysis.

About the Ipperf Option

The WAN tool **ipPerf** (referred to simply as “**ipPerf**” in this chapter) is an option of the Fabric OS **portCmd** command. This option allows you to specify the slot and port information for displaying performance statistics for a pair of ports. For this basic configuration, you can specify the IP addresses of the endpoints, target bandwidth for the path, and optional parameters such as the length of time to run the test and statistic polling interval.

Only a single **ipperf** session can be active on an FCIP GbE port at any time. Each FCIP port supports a single instance of the WAN tool-embedded client running in only sender or receiver mode. You can, however, use multiple CLI sessions to invoke simultaneous **ipPerf** sessions on different FCIP ports.

Running WAN Tool Sessions with an FCIP Tunnel Online

ipPerf sessions use different TCP ports than FCIP tunnels, so you can simultaneously run an **ipPerf** session between a pair of ports while an FCIP tunnel is online. You can, for example, revalidate the service provider Service Level Agreement (SLA) without bringing the FCIP tunnel down.

Note

For uncommitted bandwidth, any data transferred across the active tunnel will compete for the same network bandwidth as the **ipPerf** session. Unless you have a method to quiesce all storage traffic over the FCIP tunnel during **ipPerf** testing, you might experience undesirable interactions.

FCIP Port Bandwidth

Allocation of the FCIP GbE port bandwidth behaves exactly the same for **ipPerf** as for FCIP tunnels.If bandwidth is allocated for FCIP tunnels, the **ipPerf** session uses the remaining bandwidth. Since bandwidth is already reserved for the FCIP tunnels, the **ipPerf** session is not affected by any active FCIP tunnel. If no bandwidth is reserved, the **ipPerf** session competes for a share of the uncommitted bandwidth. Starting an **ipPerf** session has an impact on any active uncommitted bandwidth FCIP tunnels just like adding a new FCIP tunnel would. For example:

- Adding a committed-rate **ipPerf** session reduces the total uncommitted bandwidth shared by all the uncommitted bandwidth FCIP tunnels.
- Adding an uncommitted-bandwidth **ipPerf** session adds another flow competing for the shared uncommitted bandwidth.

The CLI and configuration system ensures that any bandwidth allocation does not result in an over commitment of the FCIP GbE port. An active FCIP tunnel cannot be forced to give up its committed buffer and bandwidth resources. Therefore, to commit a specific bandwidth to the **ipPerf** session, you must have an equivalent amount of spare capacity on the FCIP GbE port.

WAN Tool Performance Characteristics

The following table lists the end-to-end IP path performance characteristics that you can display using the **portCmd ipPerf** command and option. All four of the base **ipPerf** performance characteristics (bandwidth, loss, RTT, PMTU) are provided in the command output in Fabric OS v5.2.0 or higher.

Table 21-7 WAN Tool Performance Characteristics

Characteristic	Description
Bandwidth	Indicates the total packets and bytes sent. Bytes/second estimate are maintained as a weighted average with a 30 second sampling frequency and also as an average rate over the entire test run. The CLI output prints the bandwidth observed in the last display interval as well as the Weighted Bandwidth (WBW).BW represents what the FCIP tunnel / FC application sees for throughput rather than the Ethernet on-the-wire bytes.
Loss	Loss estimate is based on the number of TCP retransmits (assumption is that the number of spurious retransmits is minimal). Loss rate (percentage) is calculated based on the rate of retransmissions with in the last display interval.
Delay	Round Trip Time (RTT). Indicates TCP smoothed RTT and variance estimate in milliseconds.
Path MTU (PMTU)	Indicates the largest IP-layer datagram that can be transmitted over the end-to-end path without fragmentation. This value is measured in bytes and includes the IP header and payload. There is a limited support for black hole PMTU discovery. If the Jumbo PMTU (anything over 1500) does not work, ipperf will try 1500 bytes (minimum PMTU supported for FCIP tunnels). If 1500 PMTU fails, ipperf will give up. There is no support for aging. PMTU detection is not supported for active tunnels. During black hole PMTU discovery, the BW, Loss, and PMTU values printed might not be accurate.

Starting WAN Tool Analysis

Typically, you start the WAN tool before setting up a new FCIP tunnel between two sites. You can configure and use the **ipPerf** option immediately after installing the IP configuration on the FCIP port (for example, IP address, route entries). Once the basic IP addressing and IP connectivity is established between two sites, you can configure **ipPerf** with parameters similar to what will be used when the FCIP tunnel is configured.

The traffic stream generated by the WAN tool ipperf session can be used to:

- Validate a service provider Service Level Agreement (SLA) throughput, loss, and delay characteristics.
- Validate end-to-end PMTU, especially if you are trying to eliminate TCP segmentation of large Fibre Channel (FC) frames.
- Study the effects and impact FCIP tunnel traffic may have on any other applications sharing network resources.

To start an **ipPerf** session, you can use any port as long as the port (in combination with local interface) is not in use. You must run the **ipPerf** client on both the host (source mode, **-S** option) and receiver (sink mode, **-R** option). See "WAN Tool IpPerf Syntax" on page 21-29 for more information about specifying source and sink mode.

To start an ipPerf session

1. Configure the receiver test endpoint using the CP CLI.

   The syntax for invoking the receiver test endpoint using **ipPerf** for slot8, port ge0 on an FR4-18i is as follows:

   ```
   portcmd --ipperf 8/ge0 -s 192.168.255.10 -d 192.168.255.100 -R
   ```

2. Configure the sender test endpoint using a similar CP CLI.

   The syntax for invoking the sender test endpoint using **ipPerf** for slot8, port ge0 on an FR4-18i is as follows:

   ```
   portcmd --ipperf 8/ge0 -s 192.168.255.100 -d 192.168.255.10 –S
   ```

3. Use the **portCmd ipPerf** option specifying the slot and port of a switch.

The following example shows the results of the performance analysis for slot 8, port ge0:

```
ipperf to 192.168.255.10 from IP interface 192.168.255.100 on 8/0:3227
30s: BW:113.03MBps WBW(30s): 55.39MBps  Loss(%):0.0 Delay(ms):1 PMTU:1500
60s: BW:108.89MBps WBW(30s): 83.08MBps  Loss(%):0.0 Delay(ms):0 PMTU:1500
90s: BW:112.59MBps WBW(30s): 96.93MBps  Loss(%):0.0 Delay(ms):0 PMTU:1500
120s: BW:108.93MBps WBW(30s): 103.85MBps  Loss(%):0.0 Delay(ms):0 PMTU:1500
150s: BW:112.36MBps WBW(30s): 107.32MBps  Loss(%):0.0 Delay(ms):0 PMTU:1500
```

In the example above, *BW* represents Bandwidth and *Delay* represents the round-trip-time (RTT) in milliseconds. The output is printed after each display interval (if you do not specify an interval, the default is 30 seconds). In the output, data is displayed as follows:

- First column—Time elapsed since the start of the test.

- Second column—Bandwidth observed in the last display interval using the following units:
  - MBps—megabytes per second.
  - Mbps—megabits per second
  - KBps—kilobytes per second
  - Kbps—kilobits per second
  - Bps—bytes per second
  - bps—bits per second
- Third column—The 30s weighted bandwidth

WAN Tool IpPerf Syntax

When using the **portCmd ipPerf** option, you must specify the following:

- Source IP address. If the **ipPerf** is started with **–S** (source mode), the port number specified is the remote peer's port number to connect to.
- Destination IP address
- Source (transmitter) or sink (receiver). If the **ipPerf** is started with **–R** (sink mode), the port number and source IP address will be used to listen for new connections.

If the options are not specified the following defaults apply:

- Default port —3227
- Default runtime—forever
- Default display (refresh) interval—30 seconds
- Default size—1MSS

Following is the syntax for **portCmd ipPerf** to display end-to-end IP path performance statistics:

```
portCmd --ipPerf [slot]/ge# -s <source_ip> -d <destination_ip> -S | -R [-r <rate>]
[-z <size>] [-t <time>] [-i <interval>] [-p <port>]
```

where:

-s <source_ip>	Source IP address.
-d <destination_ip>	Destination IP address.
-S	Operates the WAN tool FCIP port-embedded client in the sender mode. The test endpoint will generate a traffic stream and report the end-to-end IP path characteristics from this endpoint toward the receiver endpoint. This option cannot be used with the –R option.
-R	Operates the WAN tool FCIP-port embedded client in the receiver mode. The test endpoint will accept a connection and traffic stream from the sender This option cannot be used with the -S option.
-r <*rate*>	Committed rate for the data stream in Kb/s. If specified, the traffic generator will be limited by a traffic shaper. This can be used to characterize the end-to-end IP path performance based on the data rate to be configured for a tunnel between the same endpoints. If a rate is not specified then the traffic generator will compete for uncommitted bandwidth.

-z <size>	Size in bytes for each buffer handed to the TCP layer. If a size is not specified the maximum size data buffer will be used based on the outgoing IP interface MTU. The size is the only buffer size that will be handed over to the TCP layer.
-t <time>	Total time in seconds to run the test traffic stream. If a time is not specified, the test will run continuously until the command is explicitly aborted (ctrl-C). The maximum allowed size is 1MSS. If you plan to use PMTU discovery, use the default size, otherwise the PMTU information might not be accurate.
-i <interval>	Interval in seconds between polling and printing stats. The default is 30 sec. If this time is greater than the running time (specified by –t), the stats will be printed only once at the conclusion of the test.
-p <port>	TCP port number for the listener endpoint. If a TCP port is not specified, port 3227 is used (this is next port after ports 3225 and 3226, which are used for the FCIP tunnel connections).

FCIP Tunnel Performance Characteristics

You can use the **portShow fcipTunnel** command to view the performance statistics and monitor the behavior of an online FCIP tunnel. This additional information is reported in the details of the command output.

The following example shows the use of the CLI **portCmd fcipTunnel** to report the bandwidth, loss, round trip, and path statistics for each of the TCP connections associated with active FCIP tunnel(s):

```
switch:admin06> portshow fciptunnel 8/ge0
```

To view detailed **fcipTunnel** statistics, you must specify either the **-perf** or **-params** options. The following example shows the **portCmd fcipTunnel** with the performance option to display characteristics of tunnel 0.

```
switch:admin06> portshow fciptunnel 8/ge0 all
Slot: 8 Port: ge0
----------------------------------------
        Tunnel ID 0
        Remote IP Addr 192.175.4.200
        Local IP Addr 192.175.4.100
        Remote WWN Not Configured
        Local WWN 10:00:00:60:69:e2:09:be
        Compression on
        Fastwrite off
        Committed Rate 300000 Kbps (0.300000 Gbps)
        SACK on
        Min Retransmit Time 100
        Keepalive Timeout 10
        Max Retransmissions 8
        Status : Active
        Uptime 7 minutes, 3 seconds
        QoS shaper performance stats:
            14808626616 Bytes
                39615391 Bps 30s avg, 35008573 Bps lifetime avg
            2013762456 compressed Bytes
                33208083 Bps 30s avg, 4760667 Bps lifetime avg
            7.35 compression ratio
        FC control traffic TCP connection:
            Local 192.175.4.100:4139, Remote 192.175.4.200:3225
            Performance stats:
                849 output packets
                   0 pkt/s 30s avg, 2 pkt/s lifetime avg
                173404 output Bytes
                   39 Bps 30s avg, 409 Bps lifetime avg
                0 packets lost (retransmits)
                   0.00% loss rate 30s avg
                806 input packets
                   0 pkt/s 30s avg, 1 pkt/s lifetime avg
                116736 input Bytes
                   35 Bps 30s avg, 275 Bps lifetime avg
        Data transfer TCP connection:
            Local 192.175.4.100:4140, Remote 192.175.4.200:3226
            Performance stats:
                12899612 output packets
                   34508 pkt/s 30s avg, 30495 pkt/s lifetime avg
                14499127648 output Bytes
                   38787792 Bps 30s avg, 34276897 Bps lifetime avg
                0 packets lost (retransmits)
                   0.00% loss rate 30s avg
                6495624 input packets
                   17381 pkt/s 30s avg, 15356 pkt/s lifetime avg
                207859776 input Bytes
                   556200 Bps 30s avg, 491394 Bps lifetime avg
```

The following example shows the **portCmd fcipTunnel** with the **parameters** options to display the parameters of tunnel 0:

```
switch:admin06> portshow fciptunnel 8/ge0 0
Slot: 8 Port: ge0
----------------------------------------
        Tunnel ID 0
        Remote IP Addr 192.175.4.200
        Local IP Addr 192.175.4.100
        Remote WWN Not Configured
        Local WWN 10:00:00:60:69:e2:09:be
        Compression on
        Fastwrite off
        Committed Rate 300000 Kbps (0.300000 Gbps)
        SACK on
        Min Retransmit Time 100
        Keepalive Timeout 10
        Max Retransmissions 8
        Status : Active
        Uptime 7 minutes, 3 seconds
        FC control traffic TCP connection:
            Local 192.175.4.100:4139, Remote 192.175.4.200:3225
            Runtime parameters:
                Send MSS 1456 Bytes
                Sender stats:
                    smoothed roundtrip 50 ms, variance 0
                    peer advertised window 1874944 Bytes
                    negotiated window scale (shift count) 9
                    congestion window 149649 Bytes
                    slow start threshold 1875000 Bytes
                    operational mode: slow start
                    2 packets queued: TCP sequence# MIN(2950582519)
                       MAX(2950582655) NXT(2950582655)
                    2 packets in-flight
                    Send.Unacknowledged(TCP sequence# 2950582519) recovery:
                       retransmit timeout 500 ms, duplicate ACKs 0
                       retransmits 0 (max retransmits 8)
                    loss recovery: fast retransmits 0, retransmit timeouts 0
                Receiver stats:
                    advertised window 1874944 Bytes (max 1874944)
                    negotiated window scale (shift count) 9
                    0 packets queued: TCP sequence# NXT(2101820798)
                    0 out-of-order packets queued (0 lifetime total)
                Keepalive:
                    time since last activity detected 0 s
                    idle connection probe interval 1 s
                    timeout 10 s
Data transfer TCP connection:
            Local 192.175.4.100:4140, Remote 192.175.4.200:3226
            Performance stats:
                12899612 output packets
                   34508 pkt/s 30s avg, 30495 pkt/s lifetime avg
                14499127648 output Bytes
                   38787792 Bps 30s avg, 34276897 Bps lifetime avg
                0 packets lost (retransmits)
                   0.00% loss rate 30s avg
                6495624 input packets
                   17381 pkt/s 30s avg, 15356 pkt/s lifetime avg
                207859776 input Bytes
                   556200 Bps 30s avg, 491394 Bps lifetime avg
```

The following example shows the **portShow fcipTunnel** command to display IPSec information for tunnel 0:

```
switch:admin06> portshow fciptunnel 8/ge0 3 -ipsec
Port: ge0
-----------------------------------------
        Tunnel ID 3
        Remote IP Addr 192.175.5.200
        Local IP Addr 192.175.5.100
        Remote WWN Not Configured
        Local WWN 10:00:00:05:1e:37:00:20
        Compression off
        Fastwrite on
        Tape Pipelining on
        Uncommitted bandwidth, minimum of 1000 Kbps (0.001000 Gbps)
        SACK on
        Min Retransmit Time 100
        Keepalive Timeout 80
        Max Retransmissions 9
        Status : Active
        Uptime 1 day, 23 hours, 24 minutes, 46 seconds

        IKE Policy 7
        ---------------------------------------
        Authentication Algorithm: MD5
        Encryption: 3DES
        Perfect Forward Secrecy: off
        Diffie-Hellman Group: 1
        SA Life (seconds): 200000

       IPSec Policy 7
        ---------------------------------------
        Authentication Algorithm: AES-XCBC
        Encryption: 3DES
        SA Life (seconds): 1500000

        Pre-Shared Key 1234567890123456
```

Appendix

A

Configuring the PID Format

Port identifiers (called *PIDs*) are used by the routing and zoning services in Fibre Channel fabrics to identify ports in the network. All devices in a fabric must use the same PID format, so when you add new equipment to the SAN, you might need to change the PID format on legacy equipment.

This chapter contains the following topics:

- "About PIDs and PID Binding" on page A-1
- "Summary of PID Formats" on page A-2
- "Impact of Changing the Fabric PID Format" on page A-2
- "Selecting a PID format" on page A-4
- "Evaluating the Fabric" on page A-6
- "Planning the Update Procedure" on page A-7
- "Changing to Core PID Format" on page A-9
- "Changing to Extended Edge PID Format" on page A-10
- "Converting Port Number to Area ID" on page A-12
- "Performing PID Format Changes" on page A-14
- "Swapping Port Area IDs" on page A-18

About PIDs and PID Binding

The PID is a 24-bit address built from three 8-bit fields:

- domain
- area_ID
- AL_PA

Many scenarios cause a device to receive a new PID; for example, unplugging the device from one port and plugging it into a different port as part of fabric maintenance, or changing the domain ID of a switch, which might be necessary when merging fabrics, or changing compatibility mode settings.

Some device drivers use the PID to map logical disk drives to physical Fibre Channel counterparts. Most drivers can either change PID mappings dynamically (called *dynamic PID binding*) or use the WWN of the Fibre Channel disk for mapping (called *WWN binding*).

Some older device drivers behave as if a PID uniquely identifies a device (they use *static PID binding*). These device drivers should be updated, if possible, to use WWN or dynamic PID binding instead, because static PID binding creates problems in many routine maintenance scenarios. Fortunately, very few device drivers still behave this way. Many current device drivers enable you to select static PID binding as well as WWN binding. You should only select static binding if there is a compelling reason, and only after you have evaluated the impact of doing so.

Summary of PID Formats

SilkWorm switches employ these types of PID formats:

- VC encoded —This is the format defined by the SilkWorm 1000 series. Connections to these switches are not supported in Fabric OS v4.0.0 and later.
- Native —Introduced with the SilkWorm 2000 series, this format supports up to 16 ports per switch.
- Core—The default for SilkWorm 200E, 3250, 3850, 3900, 4100, 4900 and 7500 switches and SilkWorm 24000, and 48000 directors, this is the recommended format for Brocade switches and fabrics. It uses the entire 8-bit address space and directly uses the port number as the area_ID. It supports up to 256 ports per switch.
- Extended edge —This format generates the same PID for a port on switches with 16 ports or less as would native PID format, but it also supports up to 256 ports per domain. It should be used only in cases where you cannot upgrade devices to dynamic PID binding and you absolutely cannot reboot your servers.

  Extended edge PID is supported in Fabric OS v2.6.2 and later, v3.1.2 and later, and v4.2.0 and later.

In addition to the PID formats list here, Interoperability mode supports additional PID formats that are not discussed in this guide.

Impact of Changing the Fabric PID Format

If your fabric contains switches that use Native PID, it is recommended that you change the format to Core PID before you add the new, higher port count switches and directors. Also, it is recommended that you use Core PID when upgrading the Fabric OS version on 2000 and 3000 series switches.

Depending on your situation, the PID change might or might not entail fabric downtime:

- If you are running dual-fabrics with multipathing software, you can update one fabric at a time without disrupting traffic. Move all traffic onto one fabric in the SAN and update the other fabric. Then move the traffic onto the updated fabric, and update the final fabric.
- Without dual-fabrics, stopping traffic is highly recommended. This is the case for many routine maintenance situations, so dual-fabrics are always recommended for uptime-sensitive environments. If your fabric contains devices that employ static PID binding, or you do not have dual-fabrics, you must schedule downtime for the SAN to change the PID format.

The following sections describe various impacts of PID format changes in greater detail.

Host Reboots

In some Fibre Channel SAN environments, storage devices and host servers are bound to the host operating system by their PIDs (called their *Fibre Channel addresses*). In these environments, the hosts and target HBAs in a SAN need to know the full 24-bit PIDs of the hosts and targets they are communicating with, but they do not care how the PIDs are determined. But, if a storage device PID is changed, the host must reestablish a new binding, which requires the host to be rebooted.

With the introduction of the SilkWorm 200E, 3016, 3250, 3850, 3900, 4012, and 4100 switches and the SilkWorm 12000, 24000, and 48000 directors, the Native PID format used in earlier switches was supplemented with the Core PID format, which is capable of addressing higher port counts. Changing from Native PID format to Core PID format changes the PID, which requires hosts that use port binding to be rebooted.

Static PID Mapping Errors

If you can avoid using drivers that employ static PID binding, you should do so.

With the WWN or dynamic PID binding most typically used with drivers, changing the device's PID does not affect the PID mapping. However, before updating the PID format, it is necessary to determine whether any devices in the SAN use static PID binding.

For those few drivers that do use static PID binding, changing the PID format breaks the mapping, which must be fixed either by rebooting the host or by using a manual update procedure on the host.

To correct mapping errors caused by static PID binding, see the following sections:

- See "Evaluating the Fabric" on page A-6 for details on finding devices that use static PID binding. Then see "Online Update" on page A-8 or "Offline Update" on page A-8 for recommendations.
- See "Converting Port Number to Area ID" on page A-12 for instructions.

Changes to Configuration Data

Table A-1 lists various combinations of before-and-after PID formats, and indicates whether the configuration is affected.

Caution
After changing the fabric PID format, if the change invalidates the configuration data (see Table A-1 to determine this), do not download old (pre-PID format change) configuration files to any switch on the fabric.

Table A-1 Effects of PID Format Changes on Configurations

PID Format Before Change	PID Format After Change	Configuration Effect?
Native	Extended Edge	No impact
Extended Edge	Native	No impact
Native	Core	You must: • Reenable zoning, if there is an active zone set and it uses port zones. • You do not need to reconfigure Destination ID (DID) if using: - Performance monitoring - The configure command to change the PID format - The performance monitor database - The zoning database - Security DCC database (if secure mode is enabled) The DID will be converted by Fabric OS automatically.
Core	Native	
Extended Edge	Core	
Core	Extended Edge	

After changing the fabric PID format and verifying correct fabric operation, resave configuration data by running the **configUpload** command.

Before downgrading firmware, change the PID back to supported PIDs such as Core PID. If the database is automatically converted, save the converted database, and then download the older OS.

Selecting a PID format

All switches in a fabric must use the same PID format, so if you add a switch that uses a different PID format to a fabric, the switch will segment from the fabric. The format you select for your fabric depends on the mix of switches in the fabric, and to an extent on the specific releases of Fabric OS in use (for example, Extended Edge PID format is only available in Fabric OS v2.6.2 and later, Fabric OS v3.1.2 and later, and Fabric OS v4.2.0 and later).

If you are building a new fabric with switches running various Fabric OS versions, use Core PID format to simplify port-to-area_ID mapping.

Note

Switches that are queried using outside calls should be configured using PID 1 (core PID) to ensure that the correct port numbering is used in other management applications.

Table A-2 shows various combinations of existing fabrics, new switches added to those fabrics, and the recommended PID format for that combination. The criteria for the recommendations are first to eliminate host reboots, and second to minimize the need for a host reboot in the future.

Table A-2 PID Format Recommendations For Adding New Switches

Existing Fabric OS Versions; PID Format	Switch to be Added	Recommendations (in Order of Preference)
v2.6.2 and later/v3.1.2 and later; Native PID	v2.6.2 and later/v3.1.2 and later	1. Use Native PID format for new switch Host reboot is not required. 2. Convert existing fabric to Core PID format, upgrading the version of Fabric OS, if necessary. Set Core PID format for new switch. Host reboot *is* required. 3. If devices are bound statically and it is not possible to reboot, convert existing fabric to Extended Edge PID format, upgrading the version of Fabric OS, if necessary. Use Extended Edge PID format for new switch Host reboot is not required.
	v4.2.0 and later	1. Convert existing fabric to Core PID format, upgrading the version of Fabric OS, if necessary. Set Core PID format for new switch. Host reboot *is* required. 2. If devices are bound statically and it is not possible to reboot, convert existing fabric to Extended Edge PID format, upgrading the version of Fabric OS, if necessary. Use Extended Edge PID format for new switch Host reboot is not required.
v2.6.2 and later/v3.1.2 and later/ v4.2.0 and later; Core PID	v2.6.2 and later/v3.1.2 and later/v4.2.0 and later	Use Core PID for new switch Host reboot is not required.
v2.6.2 and later/v3.1.2 and later/ v4.2.0 and later; Extended Edge PID	v2.6.2 and later/v3.1.2 and later/v4.2.0 and later	Use Extended Edge PID for new switch Host reboot is not required.

Evaluating the Fabric

If there is the possibility that your fabric contains host devices with static PID bindings, you should evaluate the fabric to:

- Find any devices that bind to PIDs
- Determine how each device driver will respond to the PID format change
- Determine how any multipathing software will respond to a fabric service interruption

If current details about the SAN are already available, it might be possible to skip the Data Collection step. If not, it is necessary to collect information about each device in the SAN. Any type of device might be able to bind by PID; each device should be evaluated before attempting an online update. This information has broad applicability, because PID-bound devices are not able to seamlessly perform in many routine maintenance or failure scenarios.

1. Collect device, software, hardware, and configuration data.

   The following is a non-comprehensive list of information to collect:

   - HBA driver versions
   - Fabric OS versions
   - RAID array microcode versions
   - SCSI bridge code versions
   - JBOD drive firmware versions
   - Multipathing software versions
   - HBA time-out values
   - Multipathing software timeout values
   - Kernel timeout values
   - Configuration of switch

2. Make a list of manually configurable PID drivers.

   Some device drivers do not automatically bind by PID, but allow the operator to manually create a PID binding. For example, persistent binding of PIDs to logical drives might be done in many HBA drivers. Make a list of all devices that are configured this way. If manual PID binding is in use, consider changing to WWN binding.

   The following are some of the device types that might be manually configured to bind by PID:

   - HBA drivers (persistent binding)
   - RAID arrays (LUN access control)
   - SCSI bridges (LUN mapping)

3. Analyze data.

   After you have determined the code versions of each device on the fabric, they must be evaluated to find out if any automatically bind by PID. It might be easiest to work with the support providers of these devices to get this information. If this is not possible, you might need to perform empirical testing.

   Binding by PID can create management difficulties in a number of scenarios. It is recommended that you not use drivers that bind by PID. If the current drivers do bind by PID, upgrade to WWN-binding drivers if possible.

The drivers shipping by default with HP/UX and AIX at the time of this writing still bind by PID, and so detailed procedures are provided for these operating systems in this chapter. Similar procedures can be developed for other operating systems that run HBA drivers that bind by PID.

There is no inherent PID binding problem with either AIX or HP/UX. It is the HBA drivers shipping with these operating systems that bind by PID. Both operating systems are expected to release HBA drivers that bind by WWN, and these drivers might already be available through some support channels. Work with the appropriate support provider to find out about driver availability.

It is also important to understand how multipathing software reacts when one of the two fabrics is taken offline. If the time-outs are set correctly, the failover between fabrics should be transparent to the users.

You should use the multipathing software to manually fail a path before starting maintenance on that fabric.

4. Perform empirical testing.

   Empirical testing might be required for some devices, to determine whether they bind by PID. If you are not sure about a device, work with the support provider to create a test environment.

   Create as close a match as practical between the test environment and the production environment, and perform an update using the procedure in "Online Update" on page A-8.

   Devices that bind by PID are unable to adapt to the new format, and one of three approaches must be taken with them:

   - A plan can be created for working around the device driver's limitations in such a way as to allow an online update. See the Detailed Procedures section for examples of how this could be done.
   - The device can be upgraded to drivers that do not bind by PID.
   - Downtime can be scheduled to reset the device during the core PID update process, which generally allows the mapping to be rebuilt.

   If either of the first two options are used, the procedures should again be validated in the test environment.

   Determine the behavior of multipathing software, including but not limited to:

   - HBA time-out values
   - Multipathing software time-out values
   - Kernel time-out values

Planning the Update Procedure

Whether it is best to perform an offline or online update depends on the uptime requirements of the site.

- An offline update that all devices attached to the fabric be offline.
- With careful planning, it should be safe to update the core PID format parameter in a live, production environment. This requires dual fabrics with multipathing software. Avoid running backups during the update process, as tape drives tend to be very sensitive to I/O interruption. The online update process is only intended for use only in uptime-critical dual-fabric environments, with multipathing software (high-uptime environments should always use a redundant fabric SAN architecture). Schedule a time for the update when the least critical traffic is running.

All switches running any version of Fabric OS 3.1.2 and later or 4.2.0 and later are shipped with the Core Switch PID Format enabled, so it is not necessary to perform the PID format change on these switches.

Migrating from manual PID binding (such as persistent binding on an HBA) to manual WWN binding and upgrading drivers to versions that do not bind by PID can often be done before setting the core PID format. This reduces the number of variables in the update process.

Online Update

The following steps are intended to provide SAN administrators a starting point for creating site-specific procedures.

1. Back up all data and verify backups.
2. Verify that the multipathing software can automatically switchover between fabrics seamlessly. If there is doubt, use the software's administrative tools to manually disassociate or mark offline all storage devices on the first fabric to be updated.
3. Verify that I/O continues over the other fabric.
4. Disable all switches in the fabric to be updated, one switch at a time, and verify that I/O continues over the other fabric after each switch disable.
5. Change the PID format on each switch in the fabric.
6. Reenable the switches in the updated fabric one at a time. In a core/edge network, enable the core switches first.
7. After the fabric has reconverged, use the **cfgEnable** command to update zoning.
8. Update their bindings for any devices manually bound by PID. This might involve changing them to the new PIDs, or preferably changing to WWN binding.

   For any devices automatically bound by PID, two options exist:

   a. Execute a custom procedure to rebuild its device tree online. Examples are provided in the "Converting Port Number to Area ID" on page A-12 section of this chapter.

   b. Reboot the device to rebuild the device tree. Some operating systems require a special command to do this, for example "boot –r" in Solaris.
9. For devices that do not bind by PID or have had their PID binding updated, mark online or reassociate the disk devices with the multipathing software and resume I/O over the updated fabric.
10. Repeat with the other fabric(s).

Offline Update

The following steps are intended to provide SAN administrators a starting point for creating site-specific procedures.

1. Schedule an outage for all devices attached to the fabric.
2. Back up all data and verify backups.

3. Shut down all hosts and storage devices attached to the fabric.
4. Disable all switches in the fabric.
5. Change the PID format on each switch in the fabric.
6. Reenable the switches in the updated fabric one at a time. In a core/edge network, enable the core switches first.
7. After the fabric has reconverged, use the **cfgEnable** command to update zoning.
8. Bring the devices online in the order appropriate to the SAN. This usually involves starting up the storage arrays first, and the hosts last.
9. For any devices manually bound by PID, bring the device back online, but do not start applications. Update their bindings and reboot again if necessary. This might involve changing them to the new PIDs, or might (preferably) involve changing to WWN binding.
10. For any devices automatically bound by PID, reboot the device to rebuild the device tree (some operating systems require a special command to do this, such as “boot –r” in Solaris).
11. For devices that do not bind by PID or have had their PID binding updated, bring them back up and resume I/O.
12. Verify that all I/O has resumed correctly.

Hybrid Update

It is possible to combine the online and offline methods for fabrics where only a few devices bind by PID. Because any hybrid procedure is extremely customized, it is necessary to work closely with the SAN service provider in these cases.

Changing to Core PID Format

In Fabric OS release v4.2.0 and later, Native PID format is not supported; the default format is the Core PID format.

In Fabric OS v3.1.2 and later, Core PID format is the default configuration.

In Fabric OS v2.6.2 and later, Native PID format is the default configuration.

Although the PID format is listed in the configuration file, do not edit the file to change the setting there. Instead, use the CLI **configure** command. When you use the **configure** command, switch databases that contain PID-sensitive information are automatically updated. If you change the setting in the configuration file and then download the edited file, the PID format will be changed, but the database entries will not, and so they will be incorrect.

The following maps the PID format names to the switch names used in the management interfaces.

- Native PID—Switch PID address mode 0
- Core PID—Switch PID address mode
- Extended edge PID—Switch PID address mode 2

Before changing the PID format, determine if host reboots will be necessary. The section "Host Reboots" on page A-3 summarizes the situations that might require a reboot.

```
switch:admin> switchdisable
switch:admin> configure
Configure...

  Fabric parameters (yes, y, no, n): [no] y

    Domain: (1..239) [1]
    BB credit: (1..27) [16]
    R_A_TOV: (4000..120000) [10000]
    E_D_TOV: (1000..5000) [2000]
    WAN_TOV: (1000..120000) [0]
    Data field size: (256..2112) [2112]
    Sequence Level Switching: (0..1) [0]
    Disable Device Probing: (0..1) [0]
    Suppress Class F Traffic: (0..1) [0]
    SYNC IO mode: (0..1) [0]
    Switch PID Address Mode: (0..2) [1] < Set mode number here.
    Per-frame Route Priority: (0..1) [0]
    Long Distance Fabric: (0..1) [0]
```

Changing to Extended Edge PID Format

In rare cases, you might be affected by the presence in the fabric of drivers that rely on static binding to the dynamically assigned PID; for example, you might be installing a switch running Fabric OS v4.2.0 into a fabric consisting solely of Fabric OS v2.6.2 and later/v3.1.2 and later switches. In these cases, if you absolutely cannot reboot the affected servers when you upgrade your switches, you can choose Extended Edge PID format. It uses the same PID mapping for the first 16 ports and can support switches and directors with higher port counts. However, because Extended Edge format only supports 128 ports per domain, its use can lead to port addressing issues in directors.

Use the following procedure only if your fabric contains devices that are bound statically and you cannot reboot the host.

1. Determine if the current switch firmware versions meet the minimum supported version levels.

   The list below provides the earliest Fabric OS version levels that support Extended Edge PID format. Use this information to determine if you need to upgrade the firmware in the switches in your fabric before you change the PID format:

   - Series 2000—v2.6.2
   - SilkWorm 3200 and 3800—v3.1.2
   - SilkWorm 3016, 3250, 3850, 3900, 12000, 24000—v4.2.0
   - SilkWorm 4012—v5.0.0
   - SilkWorm 200E, 48000—v5.0.1
   - Silkworm 4900, 7500—v5.1.0

2. Update switch firmware as necessary.

   a. Use the **fabricShow** command to verify the total number of switches in the fabric.

   b. Download the correct firmware version to each switch as necessary.

c. Reboot all switches.

d. Verify that the switches form a single fabric and that all domain IDs do not change after forming the fabric.

e. Verify that the number of switches is the same.

3. Disable the switch by entering the **switchDisable** command.

4. Change the switch configuration in the fabric to Extended Edge PID format.

   a. Configure Extended Edge PID (Format 2) on each switch. (See Figure A-1 for a sample configure command on a SilkWorm switch running Fabric OS v3.1.2 and later and see Figure A-1 for a sample configure command on a SilkWorm switch running Fabric OS 4.2.0 and later.)

   b. Run the **switchEnable** command all switches.

   c. Verify that all the switches form a fabric.

   d. Use the **switchShow** command to verify the interswitch links (ISLs) are correct and the device links are correct.

   e. Use the **fabricShow** command to verify that the number of switches are the same as those when starting this procedure.

   f. Use the **nsAllShow** command to verify the total number of devices is the same as those when starting this procedure.

   g. For dual fabrics, repeat steps 1 through 5 for the other fabric.

The following example shows the **configure** command on a switch running Fabric OS v3.1.2:

```
Configure...

Fabric parameters (yes, y, no, n): [no] yes

    Domain: (1..239) [217]
    BB credit: (1..27) [16]
    R_A_TOV: (4000..120000) [10000]
    E_D_TOV: (1000..5000) [2000]
    Data field size: (256..2112) [2112]
    Sequence Level Switching: (0..1) [0]
    Disable Device Probing: (0..1) [0]
    Suppress Class F Traffic: (0..1) [0]
    SYNC IO mode: (0..10 [0]
    Switch PID Format : (0..2) [0] 2
    Per-frame Route Priority: (0..1) [0]
    Long Distance Fabric: (0..1) [0]

   Virtual Channel parameters (yes, y, no, n): [no] ^D
 Committing configuration...done.
 0x102fd500 (tshell): Apr 15 16:53:31
    WARNING CONFIG-PIDCHANGE_DISPLACE, 3, Switch PID format changed to Extended Edge
PID Format
```

The following example shows the **configure** command on a switch running Fabric OS 5.0.1

```
configure

Configure...

  Fabric parameters (yes, y, no, n): [no] y

    Domain: (1..239) [11]
    R_A_TOV: (4000..120000) [10000]
    E_D_TOV: (1000..5000) [2000]
    WAN_TOV: (0..30000) [0]
    MAX_HOPS: (7..19) [7]
    Data field size: (256..2112) [2112]
    Sequence Level Switching: (0..1) [0]
    Disable Device Probing: (0..1) [0]
    Suppress Class F Traffic: (0..1) [0]
    Switch PID Format: (1..2) [1] 2
    Per-frame Route Priority: (0..1) [0]
    Long Distance Fabric: (0..1) [0]
    BB credit: (1..27) [16]

  Insistent Domain ID Mode (yes, y, no, n): [no]
  Virtual Channel parameters (yes, y, no, n): [no]
  Zoning Operation parameters (yes, y, no, n): [no]
  RSCN Transmission Mode (yes, y, no, n): [no]
  Arbitrated Loop parameters (yes, y, no, n): [no]
  System services (yes, y, no, n): [no]
  Portlog events enable (yes, y, no, n): [no]
  ssl attributes (yes, y, no, n): [no]
  http attributes (yes, y, no, n): [no]
  snmp attributes (yes, y, no, n): [no]
  rpcd attributes (yes, y, no, n): [no]
  cfgload attributes (yes, y, no, n): [no]
  webtools attributes (yes, y, no, n): [no]

Switch PID format has changed to PID Format 2 ('Extended Edge PID')
```

Converting Port Number to Area ID

Except for the following cases, the area ID is equal to the port number:

- When you perform a port swap operation
- When you enable Extended Edge (also known as “displaced PID”) PID on the director

If you are using Extended Edge PID format (for example, the SilkWorm 48000 with configuration option 5) and would like to map the output of the port number to the area ID, use the following formula (for ports 0–127):

```
a = (p + 16) % 128
```

where:

a = area, = port number, *%* = modulus (or remainder)

$0 \leq p < 128$

When the port number is greater than or equal to 128, the area ID and port number are the same. Figure A-1 shows a SilkWorm 48000 with Extended Edge PID.

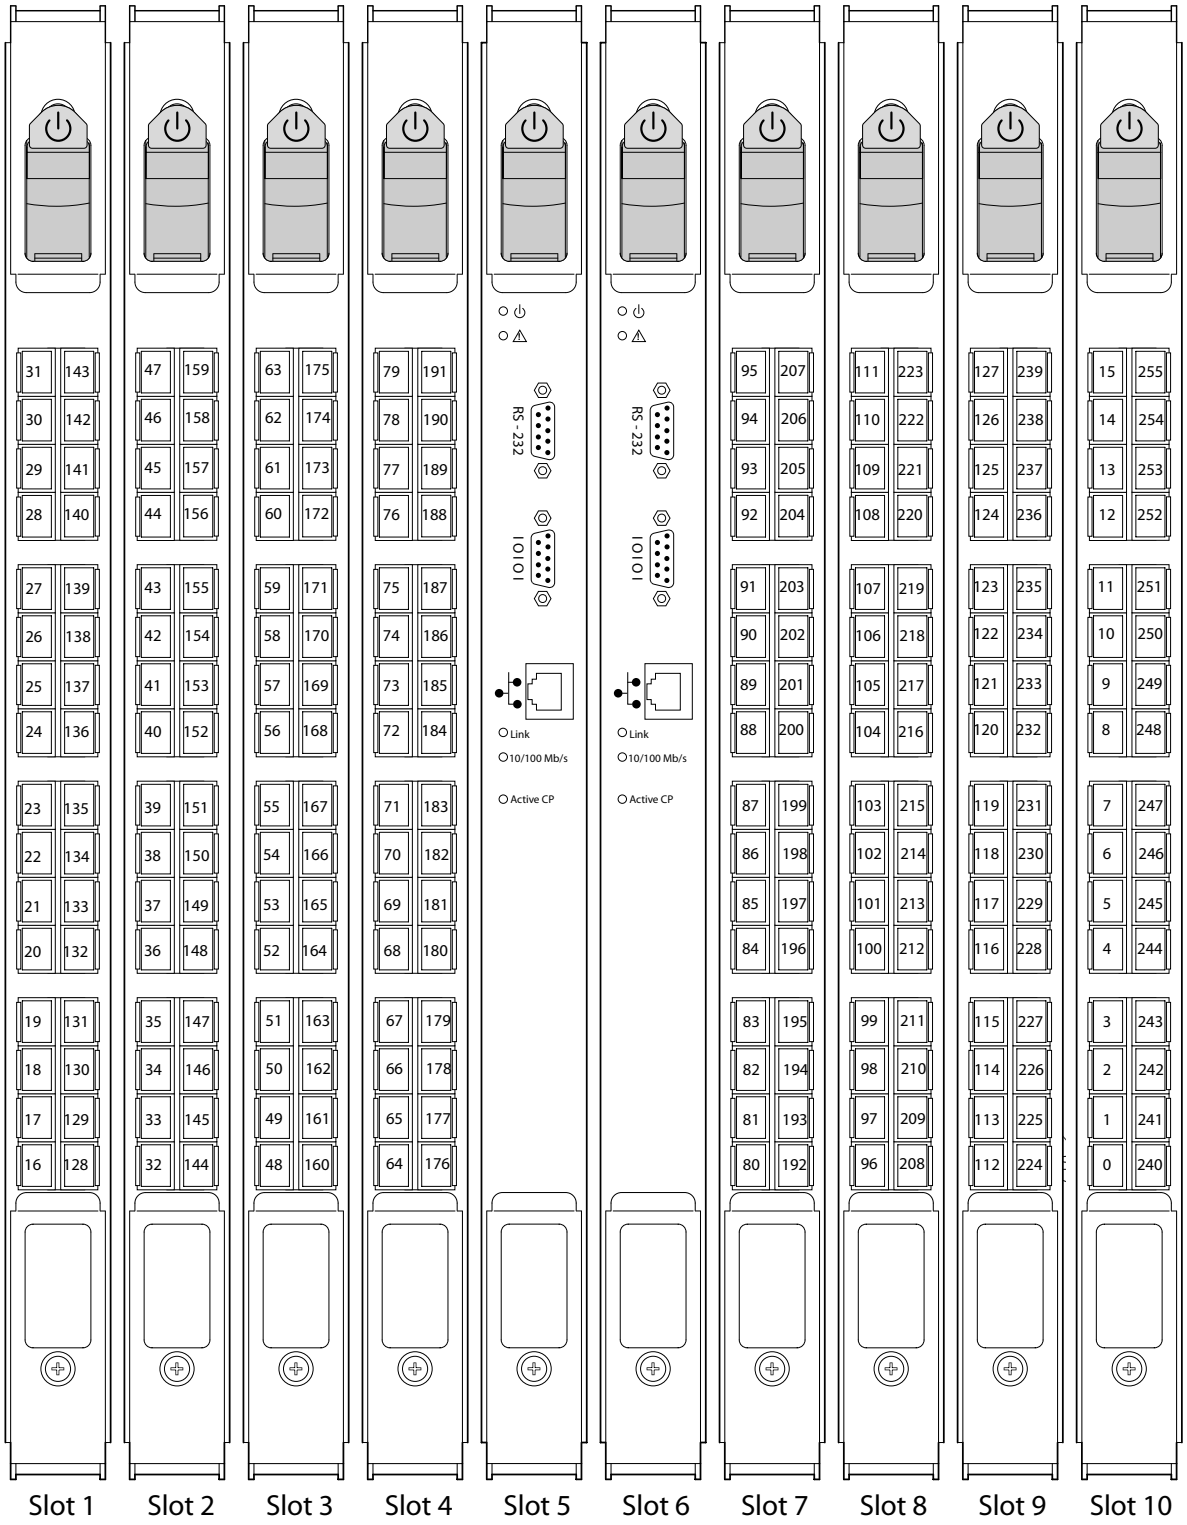

Figure A-1 SilkWorm 48000 With Extended Edge PID

Performing PID Format Changes

There are several routine maintenance procedures which might result in a device receiving a new PID. Examples include, but are not limited to:

- Changing compatibility mode settings
- Changing switch domain IDs
- Merging fabrics
- Relocating devices to new ports or new switches (that is, for Add, Move, Change type operations)
- Updating the core PID format
- Using hot spare switch ports to deal with failures

In every case where devices employ static PID binding, any such procedure becomes difficult or impossible to execute without downtime.

In some cases, device drivers allow you to specify static PID binding. In these cases, such devices must be identified and their PID binding should be changed to WWN binding.

The following sections contain a basic procedure that summarizes the steps necessary to perform PID format changes without disrupting the fabric, and special procedures for HP/UX and AIX.

Basic Procedure

This process should be executed as part of the overall online or offline update process. However, it can be implemented in a standalone manner on a non-production fabric, or a switch that has not yet joined a fabric.

1. Ensure that all switches in the fabric are running Fabric OS versions that support the addressing mode as listed below:
   - Series 2000—v2.6.2
   - SilkWorm 3200 and 3800—v3.1.2
   - SilkWorm 3016, 3250, 3850, 3900, 12000, 24000—v4.2.0
   - SilkWorm 4012—v5.0.0
   - SilkWorm 200E, 48000—v5.0.1
   - Silkworm 4900, 7500 or Silkworm 48000 with an FR4-18i blade—v5.1.0

   Note

   All switches running any version of Fabric OS 4.0.0 and later are shipped with the Core Switch PID Format enabled, so it is not necessary to perform the PID format change on these switches.

2. Telnet into one of the switches in the fabric.
3. Disable the switch by entering the **switchDisable** command.
4. Enter the **configure** command (the configure prompts display sequentially).
5. Enter **y** after the “Fabric parameters” prompt.
6. Enter **1** at the “Core Switch PID Format” prompt.

7. Complete the remaining prompts or press **Ctrl-d** to accept the remaining settings without completing all the prompts.
8. Repeat steps 2 through 7 for the remaining switches in the fabric.
9. Reenable the switch by entering the **switchEnable** command as in the following example:

```
switch:admin> switchdisable
switch:admin> configure
Configure...
    Fabric parameters (yes, y, no, n): [no] yes
    Domain: (1..239) [1]
    R_A_TOV: (4000..120000) [10000]
    E_D_TOV: (1000..5000) [2000]
    Data field size: (256..2112) [2112]
    Sequence Level Switching: (0..1) [0]
    Disable Device Probing: (0..1) [0]
    Suppress Class F Traffic: (0..1) [0]
    SYNC IO mode: (0..1) [0]
    Core Switch PID Format: (0..2) [0] 1
    Per-frame Route Priority: (0..1) [0]
    Long Distance Fabric: (0..1) [0]
    BB credit: (1..27) [16]
```

10. After all switches are updated to use the new PID format and reenabled, verify that the fabric has fully reconverged (each switch “sees” the other switches).
11. Enter **cfgEnable** [**active_zoning_config**] on one of the switches in the fabric to update zoning to use the new PID form.
    This does not change the definition of zones in the fabric, but merely causes the lowest level tables in the zoning database to be updated with the new PID format setting. It is only necessary to do this once per fabric; the zoning update automatically propagates to all switches.

    At this point, all switches in the fabric are operating in the new addressing mode.

HP/UX Procedure

This procedure is not intended to be comprehensive. It provides a starting point from which a SAN administrator could develop a site-specific procedure for a device that binds automatically by PID, and cannot be rebooted due to uptime requirements.

1. Backup all data. Verify backups.
2. If you are not using multipathing software, stop all I/O going to all volumes connected through the switch/fabric to be updated.
3. If you are not using multipathing software, unmount the volumes from their mount points using umount. The proper usage would be **umount <mount_point>**. For example:

   ```
   umount /mnt/jbod
   ```

4. If you are using multipathing software, use that software to remove one fabric’s devices from its configuration.
5. Deactivate the appropriate volume groups using vgchange. The proper usage would be **vgchange –a n <path_to_volume_group>**. For example:

   ```
   vgchange –a n /dev/jbod
   ```

6. Make a backup copy of the volume group directory using tar from within /dev. For example:

   ```
   tar –cf /tmp/jbod.tar jbod
   ```

7. Export the volume group using vgexport. The proper usage would be **vgexport –m <mapfile> <path_to_volume_group>**. For example:

   ```
   vgexport –m /tmp/jbod_map /dev/jbod
   ```

8. Connect to each switch in the fabric
9. Issue the **switchDisable** command.
10. Issue the **configure** command and change the Core Switch PID Format to 1.
11. Issue the command **cfgEnable** [**effective_zone_configuration**]. For example:

    ```
    cfgEnable my_zones
    ```

12. Issue the **switchEnable** command. Enable the core switches first, then the edges.
13. Clean the lvmtab file by using the command vgscan.
14. Change to /dev and untar the file that was tared in step 4. For example:

    ```
    tar –xf /tmp/jbod.tar
    ```

15. Import the volume groups using vgimport. The proper usage would be **vgimport –m <mapfile> <path_to_volume_group> <physical_volume_path>**. For example:

    ```
    vgimport –m /tmp/jbod_map /dev/jbod /dev/dsk/c64t8d0 /dev/dsk/c64t9d0
    ```

16. Activate the volume groups using vgchange. The proper usage would be **vgchange –a y <path_to_volume_group>**. For example:

    ```
    vgexport –a y /dev/jbod
    ```

17. If you are not using multipathing software, mount all devices again and restart I/O. For example:

    ```
    mount /mnt/jbod
    ```

18. If you are using multipathing software, reenable the affected path. The preceding steps do not “clean up” the results from ioscan. When viewing the output of ioscan, notice the that the original entry is still there, but now has a status of NO_HW.

```
# ioscan -funC disk
Class      I   H/W Path                  Driver S/W State   H/W Type       Description
----------------------------------------------------------------------------------
disk       0   0/0/1/1.2.0               adisk CLAIMED      DEVICE         SEAGATE ST39204LC
                                         /dev/dsk/clt2d0 /dev/rdsk/c1t2d0
disk       1   0/0/2/1.2.0               adisk CLAIMED      DEVICE         HP      DVD-ROM 304
                                         /dev/dsk/c3t2d0 /dev/rdsk/c3t2d0
disk       319 0/4/0/0.1.2.255.14.8.0    adisk CLAIMED      DEVICE         SEAGATE ST336605FC
                                         /dev/dsk/c64t8d0 /dev/rdsk/c64t8d0
disk       320 0/4/0/0.1.18.255.14.8.0   adisk NO_HW        DEVICE         SEAGATE ST336605FC
                                         /dev/dsk/c65t8d0 /dev/rdsk/c65t8d0
```

19. To remove the original (outdated) entry, the command rmsf (remove special file) will be needed. The proper usage for this command would be **rmsf –a –v <path_to_device>**. For example:

    ```
    rmsf –a –v /dev/dsk/c65t8d0
    ```

20. Validate that the entry has been removed by using the command **ioscan –funC disk**. In this example, the NO_HW entry is no longer listed:

```
het46 (HP-50001)> ioscan -funC disk
Class      I   H/W Path                  Driver S/W State   H/W Type       Description
-------------------------------------------------------------------------------
disk       0   0/0/1/1.2.0               adisk CLAIMED      DEVICE         SEAGATE ST39204LC
                                         /dev/dsk/clt2d0 /dev/rdsk/c1t2d0
disk       1   0/0/2/1.2.0               adisk CLAIMED      DEVICE         HP      DVD-ROM 304
                                         /dev/dsk/c3t2d0 /dev/rdsk/c3t2d0
disk       319 0/4/0/0.1.2.255.14.8.0    adisk CLAIMED      DEVICE         SEAGATE ST336605FC
                                         /dev/dsk/c64t8d0 /dev/rdsk/c64t8d0
```

21. Repeat for all fabrics.
22. Issue the **switchEnable** command. Enable the core switches first, then the edges.

AIX Procedure

This procedure is not intended to be comprehensive. It provides a starting point from which a SAN administrator can develop a site-specific procedure for a device that binds automatically by PID, and cannot be rebooted due to uptime requirements.

1. Backup all data. Verify backups.
2. If you are not using multipathing software, stop all I/O going to all volumes connected through the switch or fabric to be updated.
3. If you are not using multipathing software, varyoff the volume groups. The command usage is **varyoffvg <volume_group_name>**. For example:

   ```
   varyoffvg datavg
   ```

4. If you are not using multipathing software, unmount the volumes from their mount points using umount. The command usage is **umount <mount_point>**. For example:

   ```
   umount /mnt/jbod
   ```

5. If you are using multipathing software, use that software to remove one fabric's devices from its configuration.
6. Remove the device entries for the fabric you are migrating. For example, if the HBA for that fabric is fcs0, execute the command:

   ```
   rmdev -Rdl fcs0
   ```

7. Connect to each switch in the fabric.
8. Issue the **switchDisable** command.
9. Issue the **configure** command and change the Core Switch PID Format to 1.
10. Issue the **configEnable** [**effective_zone_configuration**] command. For example:

    ```
    configenable my_config
    ```

11. Issue the **switchEnable** command. Enable the core switches first, then the edges.
12. Rebuild the device entries for the affected fabric using the **cfgMgr** command. For example:

    ```
    cfgmgr –v
    ```

    This command might take several minutes to complete.

13. If you are not using multipathing software, vary on the disk volume groups. The proper usage would be
varyonvg <volume_group_name>. For example:

    ```
    varyonvg datavg
    ```

14. If you are not using multipathing software, mount all devices again and restart I/O. For example:

    ```
    mount /mnt/jbod
    ```

15. If you are using multipathing software, reenable the affected path.
16. Repeat for all fabrics.

Swapping Port Area IDs

If a device that uses port binding is connected to a port that fails, you can use port swapping to make another physical port use the same PID as the failed port. The device can then be plugged into the new port without the need to reboot the device.

Use the following procedure to swap the port area IDs of two physical switch ports. In order to swap port area IDs, the port swap feature must be enabled, and both switch ports must be disabled. The swapped area IDs for the two ports remain persistent across reboots, power cycles, and failovers.

SilkWorm 48000 only: You can swap only ports 0 through 15 on the FC4-48 port blade. You cannot swap ports 16 through 47.

To swap area IDs for a pair of switch ports

1. Connect to the switch and log in as admin.
2. Enable the port swap feature:

    ```
    portswapenable
    ```

3. **SilkWorm 200E, 3250, 3850, 3900, 4100, 4900 and 7500 switches**: Enter the following commands:

    ```
    portdisable port1
    portdisable port2
    ```

    SilkWorm 24000, and 48000 directors: Enter the following commands:

    ```
    portdisable slot/port1
    portdisable slot/port2
    ```

4. **SilkWorm 200E, 3250, 3850, 3900, 4100, 4900 and 7500 switches**: Enter the following command:

    ```
    portswap port1 port2
    ```

    SilkWorm 24000, and 48000 directors: Enter the following command:

    ```
    portswap slot1/port1 slot2/port2
    ```

5. Verify that the port area IDs have been swapped:

    ```
    portswapshow
    ```

    A table is shows the physical port numbers and the logical area IDs for any swapped ports.

6. Disable the port swap feature:

    ```
    portswapdisable
    ```

Appendix

B

Configuring Interoperability Mode

This appendix provides information on setting up a heterogeneous fabric that includes Brocade switches and switches from other manufacturers.

This appendix contains the following topics:

- "Vendor Switch Requirements" on page B-1
- "Brocade Switch Requirements" on page B-2
- "Supported Brocade Features" on page B-2
- "Unsupported Brocade Features" on page B-2
- "Configuration Recommendations" on page B-3
- "Configuration Restrictions" on page B-3
- "Enabling and Disabling Interoperability Mode" on page B-5

The *interoperability* mode enables Brocade switches and others to exchange interoperability parameters, allowing their fabrics to merge into one fabric with one principal switch and unique domain IDs.

The **interopMode** command must be executed on all Brocade switches in the fabric. Each Brocade switch must be rebooted after changing interoperability mode.

Switches from other manufacturers might also require the execution of one or more commands to set up interoperability mode. Refer to their documentation for details.

In a heterogeneous fabric, each Brocade switch must have interoperability mode enabled (see "Enabling and Disabling Interoperability Mode" on page B-5).

To provide maximum compatibility between switches, several features are not available in heterogeneous fabrics (see "Supported Brocade Features" on page B-2 and "Unsupported Brocade Features" on page B-2).

Vendor Switch Requirements

Interoperability has been tested on McDATA switches.

McDATA ES-3016 or ED-5000 models, or an equivalent OEM version that is plug-compatible, is required. Contact your switch supplier for the supported switches and firmware versions.

To determine whether or not a mixed-vendor SAN is supported, work with the switch provider to determine if your SAN design is valid. Important variables that determine the supportability of a particular mixed vendor SAN include the number of switches, version of Fabric OS, the topology, number of ISLs, number of connected devices, and hop count.

Brocade Switch Requirements

Following are the SilkWorm software requirements:

- 200E model running Fabric OS v5.0.1 or later.
- 2000-series models must be running Fabric OS v2.6.0 or later.
- 3000-series models must be running Fabric OS v3.1.0 or later.
- 3016, 3250, 3850, 3900, 12000, and 24000 models must be running 4.2.0 or later.
- 4100 model must be running Fabric OS v4.4.0 or later.
- 4012 model must be running Fabric OS v5.0.0 or later.
- 4900 model must be running Fabric OS v5.1.0 or later.
- 7500 model must be running Fabric OS v5.1.0 or later.
- 48000 model must be running Fabric OS v5.0.1 or later. Interoperability mode is not supported on the FC4-48 port blade.
- A zoning license and a fabric license must be installed on each Brocade switch.

Supported Brocade Features

The following features are supported on Brocade switches in interoperability mode:

- Brocade Fabric Watch
- Brocade Fabric Access API functions

  Accessible from Brocade switches only, but switch information for non-Brocade switches is reported. The object information and zoning actions are configurable from the API.

- Brocade translative mode

  Registers private storage target devices into the fabric, it can be used in a heterogeneous fabric if the devices are connected directly to Brocade switches. The devices will be accessible from any port on the fabric.

Unsupported Brocade Features

In a heterogeneous fabric, the following Brocade optional features are not supported and cannot be installed on any Brocade switch in the fabric:

- Extended Edge PID format
- Quickloop and QuickLoop Zoning
- Secure Fabric OS
- Timer Server function
- Open E_Port
- Broadcast Zoning
- Management Server Service and FDMI

- QuickLoop Fabric Assist
- Remote Switch
- Extended Fabrics
- Trunking
- Alias Server
- Platform Service
- Virtual Channels
- FCIP
- FICON

Configuration Recommendations

The following is recommended when configuring an interoperable fabric:

- Avoid domain ID conflicts before fabric reconfiguration. Every switch in the fabric must have a unique domain ID.
- When you are configuring multiple switches, you should wait for a fabric reconfiguration after adding or removing each switch.

Configuration Restrictions

In interoperable fabrics, the following restrictions apply:

- Do not use Extended Edge PID mode.
- There is an architecture maximum of 31 switches.
- Domain IDs must be in the 97 to 127 value range for successful connection to McData switches. The firmware automatically assigns a valid domain ID, if necessary, when the **interopmode** command is enabled on the switch.
- The **fabricShow** command only shows the WWN and domain ID for McDATA switches. No IP address or switch name information is provided. Brocade switches show all parameters.
- If you are managing zoning from Brocade switches, then all Brocade switches must have at least one direct connection to another Brocade switch. For example, you cannot have a McDATA switch between two Brocade switches if you are managing zoning from the Brocade switches.
- LC IBM GBICs are not supported if they are connected to a McData ISL.
- When a Brocade switch gets a new domain ID assigned through a fabric reconfiguration, the new domain ID is written to nonvolatile memory and the old domain ID value is overwritten. When a McDATA switch gets a new domain ID assigned through a fabric reconfiguration, it keeps the original domain ID in nonvolatile memory.

In this scenario, when the domain ID of both a McDATA switch and a Brocade switch are changed via fabric reconfiguration, on the next and subsequent fabric reconfiguration(s), the Brocade switch attempts to use the new ID (from the nonvolatile memory) while McDATA attempts to use its old ID (from the nonvolatile memory). This situation might cause a domain ID overlap to occur during multiple fabric reconfigurations. Domain ID overlap is not supported for Brocade/McDATA interoperability.

- Between Brocade switches, you can connect more than one ISL when in interoperability mode.

Zoning Restrictions

The following restrictions apply to zoning in interoperable fabrics:

- When interoperability mode is in effect, the space available for the zoning database is about half the usual size. The maximum zoning database size in interoperability mode is 1:1.8 of the native mode zoning database size
- Only zoning by port WWN is allowed; you must use the port WWN of the device, such as 10:00:00:00:c9:28:c7:c6.
- Zone members specified by node WWN are ignored.
- Zone configurations that use either physical port numbers or port IDs are not supported in interoperability mode. Zoning using port numbers uses the actual physical port numbers on the switch; for example slot 1, port 5.
- When a zoning configuration is not in effect, by default all ports are isolated and traffic is not permitted. This is unlike Brocade behavior where *Interoperability* mode is off (and all data traffic is enabled). If using default zoning, no device can communicate with any other device in the fabric if zoning has been disabled on a Brocade switch. See the section “Activating Default Zones” in Chapter 20, “Administering Advanced Zoning”.
- SilkWorm 200E, 3016, 3200, 3250, 3800, 3850, 3900, 4012, 4100, 24000, and 48000 models provide hardware enforcement of the port WWN zones only for devices attached to their ports. Devices attached to end-ports on non-Brocade switches or Brocade 2000-series switches are enforced by Name Server (soft) zoning only.
- Web Tools can be used for zone configuration as long as Brocade switches are connected directly to each other. If Web Tools is used to set up zoning, then it must be used as the only zone management method.
- Brocade switches connected to a McDATA switch receive the effective configuration when a zone merge occurs. (McDATA only has an effective zone configuration and discards the defined zone configuration when it sends merge information to the Brocade switch.) However, a zone update sends both the defined and the effective configuration to all switches.
- When a SilkWorm switch or director is reconfiguring, wait until the fabric routes are completely set up before entering zoning commands that must propagate to other switches. Use the **fabricShow** command to verify that all fabric routes are set up and all switch IP addresses and names are present. (The **fabricShow** command only shows the WWN and domain ID for switches from other manufacturers.)
- The maximum number of items that can be stored in the zoning configuration database depends on the switches in the fabric, whether or not interoperability mode is enabled, and the number of bytes required for each item.

You can use the **cfgSize** command to check both the maximum available size and the currently saved size. If you believe you are approaching the maximum, you can save a partially completed zoning configuration and use the cfgSize command to determine the remaining space.

Zone Name Restrictions

The name field must contain the ASCII characters that actually specify the name, not including any required fill bytes. Names must follow these rules:

- Length must be between 1 and 64 characters.
- All characters must be 7-bit ASCII.
- The first character must be a letter, which can be either upper case (A-Z) or lower case (a-z).
- Any character other than the first character must be lower case (a-z), upper case (A-Z), a single-digit number (0-9), dash (-), or underscore (_).

Enabling and Disabling Interoperability Mode

Before enabling interoperability mode, inspect the individual fabrics for compatibility.

- Make sure that zones meet the zone criteria and restrictions listed in "Zoning Restrictions" on page B-4).
- Use the **msplMgmtDeactivate** command to disable any platform management functions. (See the *Fabric OS Command Reference* for additional command information).

To enable interoperability mode

1. Verify that you have implemented all the Brocade prerequisites necessary to enable interoperability mode on the fabric (see "Configuration Recommendations" and "Configuration Restrictions" on page B-3.
2. Connect to the switch and log in as admin.
3. Enter the **switchDisable** command to disable the switch.
4. Use the **configure** command to set the domain ID to a number in the range from 97 to 126. (For detailed instructions, see "Working With Domain IDs" on page 2-17.)
5. Enter the **interopmode 1** command to enable interoperability. This command resets a number of parameters and enables interactive mode.

6. You *must* reboot the switch after changing the interoperability mode:

```
switch:admin> switchdisable
switch:admin> configure
Configure...
Fabric Parameters (yes, y, no, n): [no] y
Domain (1...239): [1] 97
...

switch:admin> interopmode 1
Please make sure the MS Platform service is disabled.
Please run "msPlMgmtDeactivate" command to disable the Platform service.
The switch effective configuration will be lost when the operating mode is
changed; do you want to continue? (yes, y, no, n): [no] y
Interopmode is enabled
Note: You must reboot this switch for the new change to take effect.
switch:admin>
```

7. Repeat step 2 through step 6 on all Brocade switches in the fabric.
8. Enable interoperability on each non-Brocade switch (refer to the switch documentation).
9. After enabling interoperability mode on all switches, physically connect the non-Brocade switches into the Brocade fabric, one at a time.

To disable interoperability mode

1. Connect to the switch and log in as admin.
2. Enter the **switchDisable** command to disable the switch.
3. Enter the **interopmode 0** command to disable interoperability. This command resets a number of parameters and disables interactive mode.
4. You *must* reboot the switch after changing the interoperability mode:

```
switch:admin> switchdisable
switch:admin> interopmode 0
The switch effective configuration will be lost when the operating mode is
changed; do you want to continue? (yes, y, no, n): [no] y

done.
Interopmode is disabled

Note: You must reboot this switch for the new change to take effect.
switch:admin>
```

5. Wait for a fabric reconfiguration after removing each switch.
6. Each non-Brocade switch might require the execution of a similar command to disable interoperability.
7. Repeat this procedure on all Brocade switches in the fabric.

Understanding Legacy Password Behavior

Appendix C

This appendix provides password information for early versions of Fabric OS firmware and contains the following topics:

- "Password Management Information" on page C-1
- "Password Prompting Behaviors" on page C-3
- "Password Migration During Firmware Changes" on page C-4
- "Password Recovery Options" on page C-5

Password Management Information

The following table describes the password standards and behaviors between various versions of firmware.

Table C-1 Account/Password Characteristics Matrix

Topic	v4.0.0	v4.1.0 to v4.2.0	v4.4.0 to 5.1.0
Number of default accounts on the switch	4, chassis based	**12000** - 8 for the chassis, 4 per switch **3016, 3250, 3850, 3900, 24000** - 4	**12000** - 8 for the chassis, 4 per switch **3016, 3250, 3850, 3900, 4100, 24000** - 4
Default account names	root, factory, admin, user	root, factory, admin, user. **3016:** administrative account is called USERID.	root, factory, admin, user. **3016:** administrative account is called USERID.
Account name changing feature	No	No, regardless of security mode.	**3016** - supports the **userRename** command.
Maximum and minimum amount of characters for a password	0 - 8 (Standard UNIX)	8 - 40 characters with printable ASCII	8 - 40 characters with printable ASCII **Note:** The minimum password length for v5.1.0 is configurable.

Table C-1 Account/Password Characteristics Matrix (Continued)

Topic	v4.0.0	v4.1.0 to v4.2.0	v4.4.0 to 5.1.0
Can different switch instances use a different password for the same account login level? For example, the password for admin for switch 0 can be different from password for admin for switch 1.	No	Yes for 12000 switch. n/a for all other switches.	Yes for 12000 switch. n/a for all other switches.
Does the root account use restricted shell?	No	No	No
When connecting to a factory installed switch, do you use the default passwords?	Yes	Yes	Yes
Does a user need to know the old passwords when changing passwords using the **passwd** command?	Yes, except when the root user changes another user's password. This is standard UNIX behavior; Fabric OS does not enforce any additional security.	Old password is required only when changing password for the same level user password. Changing password for lower level user does not require old password. For example, users connect as admin; old admin password is required to change the admin password. But old user password is not required to change the user password.	Old password is required only when changing password for the same level user password. Changing password for lower level user does not require old password. For example, users connect as admin; old admin password is required to change the admin password. But old user password is not required to change the user password.
Can **passwd** change higher-level passwords? For example, can admin change root password?	Yes, but will ask for the "old password" of the higher-level account (example "root").	Yes; if users connect as admin, they can change the root, factory, and admin passwords. However, if one connects as user, one can only change the user password.	Yes, if users connect as admin, they can change the root, factory and admin passwords after first entering the old password for the respective account.
Can API change passwords?	Yes, only for admin.	Yes, only for admin.	Yes, only for admin.
Can Web Tools change passwords?	No	No	No
Can SNMP change passwords?	No	No	No

Password Prompting Behaviors

The following table describes the expected password prompting behaviors of various Fabric OS versions.

Table C-2 Password Prompting Matrix

Topic	v4.0.0	v4.1.0 and later
Must *all* password prompts be completed for *any* change to take effect?	No. Partial changes of all four passwords are allowed.	No. Partial changes of all four passwords are allowed.
When does the password prompt appear?	When users connect as root, factory, or admin, the accounts with default password will be prompted for change. The accounts with non-default password will NOT be prompted.	When users connect as root, factory, or admin, the accounts with default password will be prompted for change. The accounts with non-default password will NOT be prompted.
Is a user forced to answer password prompts before getting access to the firmware?	No, users can type in **Ctrl-c** to get out of password prompting.	No, users can type in **Ctrl-c** to get out of password prompting.
Do users need to know the old root password when answering prompting?	Yes in v4.0.0 *No in v4.0.2 only	No
Are new passwords forced to be set to something different than the old passwords?	Yes	Yes
Is password prompting disabled when security mode is enabled?	Yes	Yes
Is the **passwd** command disabled until the user has answered password prompting?	No	Yes
Does password prompting reappear when passwords are changed back to default using the **passwd** command?	Yes	No
Does password prompting reappear when passwords are changed back to default using the **passwdDefault** command?	Yes	Yes

Password Migration During Firmware Changes

Table C-3 describes the expected outcome of password settings when upgrading or downgrading firmware for various Fabric OS versions.

Table C-3 Password Migration Behavior During Firmware Upgrade/Downgrade

Topic	v4.4.0 to v5.0.1	v5.0.1 to v5.1.0
Passwords used when upgrading to a newer firmware release for the first time.	Default accounts and passwords are preserved.	Default accounts and passwords are preserved.
Passwords preserved during subsequent firmware upgrades	MUA accounts created during a previous upgrade to v5.0 are restored.	MUA accounts created during a previous upgrade to v5.1.0 are restored.
Passwords used if downgrading to an older firmware for the first time	Downgrades to v4.4.0 preserve all existing default accounts, MUA accounts and passwords. MUA accounts with the switchAdmin role have the same permissions as the user role.	Downgrades to v5.0.1 preserve all existing default accounts, MUA accounts and passwords.
When downgrading to an older firmware at subsequent times, which passwords will be used?	Downgrades to v4.4.0 preserve all existing default accounts, MUA accounts and passwords. MUA accounts with the switchAdmin role have the same permissions as the user role.	Downgrades to v5.0.1 preserve all existing default accounts, MUA accounts and passwords.
When downgrading then upgrading again, what passwords will be used?	All default and MUA accounts and passwords remain unchanged.	All default and MUA accounts and passwords remain unchanged.

Password Recovery Options

The following table describes the options available when one or more types of passwords are lost.

Table C-4 Password Recovery Options

Topic	v4.0.0	v4.1.0 and later
If all the passwords are forgotten, what is the password recovery mechanism? Are these procedures non-disruptive recovery procedures?	Contact your switch service provider. A non-disruptive procedure is available.	Contact your switch service provider. A non-disruptive procedure is available.
If a user has only the root password, what is the password recovery mechanism?	Root can change any password by using the **passwd** command.	Use **passwd** command to set other passwords. Use **passwdDefault** command to set all passwords to default.
How to recover boot PROM password?	n/a	Contact your switch service provider and provide the recovery string. Refer to "Setting the Boot PROM Password" on page 3-30 for more information.
How do I recover a user, admin, or factory password?	Contact your switch service provider. Refer to "Recovering Forgotten Passwords" on page 3-35 for more information.	

Using Remote Switch

Appendix D

This appendix describes the concepts and procedures for using the Brocade Remote Switch feature and contains the following topics:

- "About Remote Switch" on page D-1
- "Remote Switch Capabilities" on page D-1
- "Using Remote Switch with a Gateway" on page D-2

About Remote Switch

The Brocade Remote Switch feature, which aids in ensuring gateway compatibility, was formerly a licensed feature. Its functionality is now available as part of the Fabric OS standard feature set through the use of the **portCfgIslMode** command, which is described in "Linking Through a Gateway" on page 2-23. For those who use Remote Switch as part of their legacy set of tools, this appendix contains a description and procedure for the feature.

Remote Switch, enables you to connect two remote Brocade switches over an IP network, enabling communication of IP or ATM protocols as well as Fibre Channel traffic.

The Brocade Remote Switch feature functions with the aid of a "bridging device" or Fibre Channel gateway. The gateway supports both a Fibre Channel physical interface and a secondary, non-Fibre Channel physical interface, such as IP, SONET, or ATM. Remote Switch functions over E_Port connections. With Remote Switch on both fabrics, the gateway accepts Fibre Channel frames from one fabric, tunnels them across the network, and passes them to the other fabric. From the viewpoint of the connected hosts and storage devices, fabrics using Remote Switch interact the same as locally connected switches.

Remote Switch Capabilities

Remote Switch provides many of the same capabilities of normal ISL links including,

- Coordinated fabric services

  The Remote Switch fabric configuration fully supports all fabric services, including distributed name service, registered state change notification, and alias service.

- Distributed management

  Management tools such as Advanced Web Tools, Fabric OS, and SNMP are available from both the local switch and the remote switch. Switch management is routed through the Fibre Channel connection; thus, no additional network connection is required between sites.

- Support for interswitch links (ISLs)

  Sites requiring redundant configurations can connect multiple E_Ports to remote sites by using multiple gateways. Standard Fabric OS routing facilities automatically maximize throughput and provide automatic failover during interruption on the WAN connection.

Using Remote Switch with a Gateway

The Remote Switch feature operates in conjunction with a gateway. The gateway provides an E_Port interface that links to the SilkWorm E_Port. After the link between the two E_Ports has been negotiated, the gateway E_Port moves to passthrough mode and passes Fibre Channel traffic from the SilkWorm E_Port to the WAN.

The gateway accepts Fibre Channel frames from one side of a Remote Switch fabric, transfers them across a WAN, and passes them to the other side of the Remote Switch fabric.

Remote Switch can be used for these types of gateway devices:

- Fibre Channel over ATM
- Fibre Channel over IP
- Fibre Channel over SONET
- Fibre Channel over DWDM

Most of these gateway devices have enough buffers to cover data transfer over a wide area network (WAN). The SilkWorm switches on each side of the gateway must have identical configurations. Only qualified SFPs should be used.

You must connect the fabrics through the gateway device, and make sure that the **configure** parameters are compatible with the gateway device.

You might be required to reconfigure the following parameters, depending on the gateway requirements:

Note

Consult your gateway vendor for supported and qualified configurations.

- R_A_TOV: Specify a Resource Allocation Timeout Value compatible with your gateway device.
- E_D_TOV: Specify a Error Detect Timeout Value compatible with your gateway device
- Data field size: Specify the maximum Fibre Channel data field reported by the fabric. Verify the maximum data field size the network-bridge can handle. Some bridges might not be able to handle a maximum data field size of 2112.
- BB credit: Specify the number of Buffer-to-Buffer credits for Nx_port devices.
- Suppress Class F Traffic: Use this parameter to disable class F traffic. Some network-bridge devices might not have a provision for handling class F frames. In this case, the transmission of class F frames must be suppressed throughout the entire Remote Switch fabric.

To set the access and reconfigure these parameters:

1. Connect to the switch and log in as admin.
2. Enter the **switchDisable** command to disable the switch.
3. Enter the **configure** command.

4. At the Fabric Parameters prompt enter **yes**.
5. Press **Enter** to scroll through the **Fabric Parameters** without changing their values, until you reach the parameter you want to modify.
6. Specify a new parameter value that is compatible with your gateway device.
7. Press **Enter** to scroll through the remainder of the configuration parameters. Make sure that the configuration changes are committed to the switch.
8. Repeat for all switches in the fabrics to be connected through a gateway device. These parameters must be identical on each switch in the fabric, and between fabrics connected through the gateway device.

This example shows how to modify the data field size and suppress class F traffic on a switch:

```
switch:admin> switchdisable
switch:admin> configure
Configure...
Fabric parameters (yes, y, no, n): [no] yes
Domain: (1..239) [3]
    R_A_TOV: (4000..120000) [10000]
    E_D_TOV: (1000..5000) [2000]
    Data field size: (256..2112) [2112] 1000
    Sequence Level Switching: (0..1) [0]
    Disable Device Probing: (0..1) [0]
    Suppress Class F Traffic: (0..1) [0] 1
    VC Encoded Address Mode: (0..1) [0]
    Per-frame Route Priority: (0..1) [0]
    Long Distance Fabric: (0..1) [0]
    BB credit: (1..16) [16]
Virtual Channel parameters (yes, y, no, n): [no]
  Zoning Operation parameters (yes, y, no, n): [no]
  RSCN Transmission Mode (yes, y, no, n): [no]
  NS Operation Parameters (yes, y, no, n): [no]
  Arbitrated Loop parameters (yes, y, no, n): [no]
  System services (yes, y, no, n): [no]
  Portlog events enable (yes, y, no, n): [no]
Committing configuration...done.
switch:admin>
```

Appendix

E

Zone Merging Scenarios

Table E-3 provides information on merging zones and the expected results.

Table E-3 Zone Merging Scenarios

Description	Switch A	Switch B	Expected Results
Switch A with a defined configuration **Switch B** does not have a defined configuration	defined: cfg1: zone1: ali1; ali2 effective: none	defined: none effective: none	Configuration from **Switch A** to propagate throughout the fabric in an inactive state, because the configuration is not enabled.
Switch A with a defined and enabled configuration **Switch B** has a defined configuration but no effective configuration	defined: cfg1 zone1: ali1; ali2 effective: cfg1:	defined: cfg1 zone1: ali1; ali2 effective: none	Configuration from **Switch A** to propagate throughout the fabric. The configuration is enabled after the merge in the fabric.
Switch A and **Switch B** have the same defined configuration. Neither have an enabled configuration.	defined: cfg1 zone1: ali1; ali2 effective: none	defined: cfg1 zone1: ali1; ali2 effective: none	No change (clean merge).
Switch A and **Switch B** have the same defined and enabled configuration.	defined: cfg1 zone1: ali1; ali2 effective: cfg1:	defined: cfg1 zone1: ali1; ali2 effective: cfg1:	No change (clean merge).
Switch A does not have a defined configuration **Switch B** with a defined configuration	defined: none effective: none	defined:cfg1 zone1: ali1; ali2 effective: none	**Switch A** will absorb the configuration from the fabric.
Switch A does not have a defined configuration **Switch B** with a defined configuration	defined: none effective: none	defined:cfg1 zone1: ali1; ali2 effective: cfg1	**Switch A** will absorb the configuration from the fabric - with cfg1 as the effective cfg.
Switch A and **Switch B** have the same defined configuration. Only **Switch B** has an enabled configuration.	defined: cfg1 zone1: ali1; ali2 effective: none	defined: cfg1 zone1: ali1; ali2 effective: cfg1	Clean merge - with cfg1 as the effective cfg.

Table E-3 Zone Merging Scenarios (Continued)

Description	Switch A	Switch B	Expected Results
Switch A and **Switch B** have different defined configurations. Neither have an enabled zone configuration.	defined: cfg2 zone2: ali3; ali4 effective: none	defined: cfg1 zone1: ali1; ali2 effective: none	Clean merge - the new cfg will be a composite of the two -- defined: cfg1 zone1: ali1; ali2 cfg2: zone2: ali3; ali4 effective: none
Switch A and **Switch B** have different defined configurations. **Switch B** has an enabled configuration.	defined: cfg2 zone2: ali3; ali4 effective: none	defined: cfg1 zone1: ali1; ali2 effective: cfg1	Clean merge - The new cfg will be a composite of the two, with cfg1 as the effective cfg.
Effective cfg mismatch	defined: cfg1 zone1: ali1; ali2 effective: cfg1 zone1: ali1; ali2	defined: cfg2 zone2: ali3; ali4 effective: cfg2 zone2: ali3; ali4	Fabric segments due to: Zone Conflict cfg mismatch
cfg content mismatch	defined: cfg1 zone1: ali1; ali2 effective: irrelevant	defined: cfg1 zone1: ali3; ali4 effective: irrelevant	Fabric segments due to: Zone Conflict content mismatch
	defined: cfg1 zone1: ali1; ali2 effective: irrelevant	defined: cfg1 zone1: ali1; ali4 effective: irrelevant	Fabric segments due to: Zone Conflict content mismatch
Same content - different effective cfg name	defined: cfg1 zone1: ali1; ali2 effective: cfg1 zone1: ali1; ali2	defined:cfg2 zone1: ali1; ali2 effective:cfg2 zone1: ali1; ali2	Fabric segments due to: Zone Conflict cfg mismatch
Same content - different zone name	defined: cfg1 zone1: ali1; ali2 effective: irrelevant	defined: cfg1 zone2: ali1; ali2 effective: irrelevant	Fabric segments due to: Zone Conflict content mismatch
Same content - different alias name	defined: cfg1 ali1: A; B effective: irrelevant	defined:cfg1:ali2: A; B effective: irrelevant	Fabric segments due to: Zone Conflict content mismatch
Same name - different types	effective: zone1: MARKETING	effective: cfg1: MARKETING	Fabric segments due to: Zone Conflict type mismatch
Same name - different types	effective: zone1: MARKETING	effective: alias1: MARKETING	Fabric segments due to: Zone Conflict type mismatch
Same name - different types	effective:cfg1: MARKETING	effective:alias1: MARKETING	Fabric segments due to: Zone Conflict type mismatch

Index

A

AAA service requests 3-17
aaaConfig command 3-6, 3-7, 3-27, 3-29
access
  active ports 2-19
  browser support 4-7
  changing account parameters 3-9
  control 4-17
  CP blade 3-22
  creating accounts 3-8
  deleting accounts 3-9
  full fabric 2-19
  IP address changes 2-2
  log in fails 2-2
  NTP 2-14
  other devices 2-22
  other switches 2-22
  password, changing 2-5
  recovering accounts 3-10
  remote access policies 3-26
  secure, HTTPS 4-7
  secure, SSL 4-7
  SNMP ACL 4-1
access methods
  configuration, CLI 1-2
  configuration, Fabric Manager 1-2
  configuration, Web Tools 1-2
accessControl command 4-17
accessing switches and fabrics 4-5
account ID 2-3
account lockout policy 3-16
  duration 3-16
  threshold 3-16
accounts
  changing parameters 3-9
  creating 3-8
  default, SilkWorm 12000 and 24000 2-4
  deleting 3-9
  displaying information 3-8
  domain defaults 2-4
  managing passwords 3-11
  password rules 3-11
  recovering 3-10
  user-defined 3-7
activating
  a switch certificate 4-11
  Admin Domains 7-16
  certificates 4-11
  POD 2-20
  policy 5-8
  policy changes 5-8
  ports on demand 2-19
  unlicensed ports 2-19
AD0 7-4
AD255 7-5
adding
  a new switch or fabric to a zone 20-26
  Admin Domain members 7-17
  alias members 20-12
  and removing FICON CUP licenses 12-17
  custom filter-based monitors 17-10
  end-to-end monitors 17-4
  filter-based monitors 17-9
  members to a zone configuration 20-20
  port mirror connection 15-31
  RADIUS configuration 3-28
  standard filter-based monitors 17-9
  switches to a zone 20-26
  zone members 20-14
ADList 3-20
admin, default, SilkWorm 12000and 24000 2-4
administrative domains
  about 7-1
  activating and deactivating 7-16
  AD0 7-4
  AD255 7-5
  adding members 7-17
  assigning users to 7-15
  configuration, displaying 7-22
  creating 7-13
  deleting 7-19
  firmware upgrade 7-11
  implementing 7-13
  interaction with Fabric OS features 7-24

logging in to 7-7
member types 7-8
removing members 7-17
renaming 7-18
switch WWN 7-9
switching context 7-23
transaction model 7-12
using 7-21
validating members 7-20
advanced performance monitoring commands 17-2
agtCfgDefault command 4-14, 4-21
agtCfgSet command 4-14, 4-20
agtCfgShow command 4-14
AIX procedure, PID A-17
AL_PA 11-33
alias
adding members 20-12
creating 20-11
deleting 20-12
removing members 20-12
analyzing connection problems 15-5
assigning a static route 10-3
assigning users to Admin Domains 7-15
authenticating users 3-1
authentication
configuring 3-6
local 3-30
auto-leveling, FR4-18i blade 8-11, 8-17

B

backbone fabric ID 11-13
backbone-to-edge routing 11-9, 11-13
backing up
configurations 6-2
and restoring configurations, FICON 12-18
basic
card management 9-7
connections 2-22
PID procedure A-14
beaconing mode 9-17
blades 9-1
blocking listeners 4-5
boot password 3-30
boot PROM password 3-30
Brocade Vendor-Specific Attribute 3-17
browser
configuring 4-11
Java support 4-7
troubleshooting certificates 4-13
buffer-limited ports 18-2

C

card management 9-7
certificates
activating 4-11
browser, configuring 4-11
choosing 4-8
corrupt 4-13
deleting 4-13
displaying 4-13
HTTPS 4-2
installing 4-10
invalid 4-13
not installed 4-13
obtaining 4-10
private key 4-8
public key 4-8
root 4-8
root, configuring 4-11, 4-12
secure telnet 4-2
security 4-2
SSH 4-2
SSL 4-2, 4-7, 4-8
switch 4-8
troubleshooting 4-13
cfgActvShow command 11-50
changing
an account password 3-13
chassis name 2-17
configuration data A-3
RADIUS configuration 3-29
RADIUS servers 3-29
SNMP MIB trap values 4-22
SNMP values 4-20
switch names 2-16
to core PID format A-9
to extended edge PID format A-10
CHAP
account policies 3-24
enabling 3-24
chassis name 2-17
chassisShow command 2-24

checking
connected switches 8-5
status 2-24
choosing
a CA 4-8
extended ISL mode 18-3
clearing
FICON management database 12-4, 12-11
monitor counters 17-15
CLI 1-2
access methods 1-2
overview 1-2
supported tools 1-2
collecting performance data 17-17
combining SilkWorm 12000 and 24000 cards in one chassis 9-17
commands 11-12
aaaConfig 3-6, 3-7
advanced performance monitoring 17-2
cfgActvShow 11-50
chassisName 2-17
chassisShow 2-24
configUpload 5-3, 6-10
defZone 11-23
fabricShow 11-43, 2-24
fcrConfigure 11-13
fcrProxyConfig 11-29
fcrProxyDevShow 11-47
fcrXlateConfig 11-29
frcResourceShow 11-33
fosConfig 11-13
haShow 2-24
interopMode 11-12
licenseAdd 2-15
licenseRemove 2-16
licenseShow 2-15, 2-16
lsanZoneShow 11-50
nsAllShow 11-43, 11-49, 2-25
nsCamShow 11-49
nsShow 2-25
passwdCfg 3-13
portCfgEXPort 11-40
portCfgLongDistance 2-22
portDisable 2-21, 2-22
portEnable 2-21
portLog 11-47
portShow 11-40
portStop 11-40
secPolicyAbort 5-10
secPolicyActivate 5-6, 5-8, 5-9, 5-10
secPolicyAdd 5-9
secPolicyDelete 5-10
secPolicyRemove 5-9
secPolicySave 5-6, 5-8, 5-10
slotShow 2-24, 11-11
supportSave 11-8
supportShow 11-8
switchDisable 2-20
switchEnable 2-20
switchName 2-17
switchShow 11-42, 2-24
userConfig 3-3
version 11-11
common problem areas 15-2
configuration
FICON environment switched point-to-point 12-3
FICON environment, cascaded 12-3
high-integrity fabric 12-7
new SilkWorm 24000 with two domains 9-14
recommendations for interoperability B-3
restrictions for interoperability B-3
save to a host 6-1
settings, FICON environment 12-6
configuring 4-15, 4-23
access control 4-17
access methods, CLI 1-2
access methods, Fabric Manager 1-2
access methods, Web Tools 1-2
authentication 3-6
browser 4-11
browser, certificates 4-11
certificates 4-7
changing RADIUS servers 3-29
CP switch, RADIUS 3-23
date and time 2-9
extended ISL 18-10
FibreAlliance MIB 4-14
FICON environment switches for 12-6
for interconnectivity 11-39
host 14-14
HTTPS access 4-7
IAS 3-24
IAS user 3-24
interfabric link 11-16
LINUX RADIUS server 3-22
NTP 2-14
private key 4-8
public key 4-8
RADIUS 3-25
RADIUS server 3-22
RADIUS, changing 3-29

root certificates 4-11, 4-12
secure file copy 4-23
security features 4-1, 5-1
security levels 4-15
server database 13-6
SNMP 4-14
SNMP traps 4-14
SSH client 4-3
SSL 4-7
SSL protocol 4-7
switch 3-27, 12-6
switch for RADIUS 3-27
switch, FICON environment 12-6
switch, RADIUS client 3-23
switch, single 12-7
syslogd 14-13
telnet interface 4-4
Windows RADIUS client 3-24
zone, rules for 20-10
connecting
multiple EX_Ports to an edge fabric 11-11
other devices 2-22
other switches 2-22
to devices 2-22
to other switches 2-22
to the CLI 2-2
connection
restrictions 3-2
serial 2-2
telnet 2-2
connection unit
state 4-22
status 4-22
connectivity
event 4-22
event description 4-22
status 4-22
connUnitEventDescr command 4-22
connUnitEventId command 4-22
connUnitEventObject command 4-22
connUnitEventTrap command 4-22
connUnitEventType command 4-22
connUnitPortState command 4-22
connUnitPortStatus command 4-22
connUnitPortStatusChange command 4-22
connUnitSensorStatus command 4-22
connUnitSensorStatusChange command 4-22
connUnitState command 4-22
connUnitStatus command 4-22
connUnitStatusChange command 4-22
conserving power 9-9
consolidating front domains 11-5
controlling access 13-2
controlling topology discovery 13-7
converting an installed SilkWorm 24000 to support
two domains 9-15
core/edge topology and ISL trunking 19-3
correcting
device login issues 15-14
I2C bus errors 15-13
link failures 15-21
marginal links 15-24
zoning setup issues 15-10
CP blade 2-15
access 3-22
CP switch
RADIUS configuration 3-23
CRC errors, displaying 17-3
creating
accounts 3-8
Admin Domains 7-13
alias 20-11
DCC policy 5-4
IPSec tunnel 21-11
policy 5-4, 5-7
SCC policy 5-7
zone configurations 20-19, 20-20
zones 20-13
CSR 4-9
customizing
chassis name 2-17
switch names 2-16

D

database, clearing in a FICON environment 12-4
date 2-9
date and time 2-9, 2-13
DCC policy 5-4
deactivating Admin Domains 7-16
default names 2-16
default zone mode 20-15

default switch names 2-16
defined AD configuration 7-12
defined zone configuration 20-6
defZone command 11-23
deleting
accounts 3-9
Admin Domains 7-19
alias 20-12
end-to-end monitors 17-8
filter-based monitors 17-11
policy 5-10
port mirror connection 15-31
RADIUS configuration 3-29
zone configurations 20-21
zones 20-14
designing fabric for trunking 19-3
deskew 11-33
deskew values, displaying 19-8
devices
connecting 2-22
proxy 11-8
DH-CHAP 11-14
DH-CHAP secret 11-15
dictionary.brocade 3-19
disabled zone configuration 20-6
disabling
ports 2-21
interoperability mode B-6
RADIUS configuration 3-28
switches 2-20
disabling and enabling
ports 2-21
switches 2-17
cards 9-8
displaying
additional help topics 1-4
command help 1-3
configuration settings 6-1
CRC error count 17-3
end-to-end mask 17-7
fmsmode setting 12-14
information 12-10
mode register bit settings 12-15
monitor counters 17-13
node identification data, FICON environments 12-10
RADIUS configuration 3-27
registered listeners for link incidents, FICON environment 12-10
trunking information 19-8
displaying and clearing the CRC error count 17-3
displaying and deleting certificates 4-13
domain
default accounts 2-4
default roles 2-4
domain ID 2-16
FICON environment 2-18
insistent 12-1
domain, phantom 11-10
downgrade and port mirroring 15-30
downloading
configurations across a fabric 6-9
firmware 8-2

E

edge-to-edge routing 11-13
EFCM 11-43, 11-48
effective
AD configuration 7-12
zone configuration 20-6
enabling
CHAP 3-24
interoperability mode B-5
port 2-21
switch 2-20
enabling and disabling
FICON management server mode 12-13
interoperability mode B-5
ISL trunking 19-6
local authentication 3-30
platform services 13-1
encryption 4-7
end-to-end monitoring 17-3
end-to-end monitors
adding 17-4
deleting 17-8
displaying the mask 17-7
restoring configuration 17-16
saving configuration 17-16
setting a mask 17-6
ensuring network security 4-3
evaluating the fabric A-6

event
connectivity 4-22
date and time 2-9
description, connectivity 4-22
sensor 4-22
EX_Ports 11-35, 11-42, 15-27
exchange-based routing 10-2
Extended Fabrics 18-1
extended ISL
choosing mode 18-3
configuring 18-10
extended link buffer allocation 18-1

F

fabric, high integrity 12-3
fabrics
designing for trunking 19-3
full access 2-19
access 4-5
connectivity 2-24
considerations 19-3
Fabric Manager, access methods 1-2
Fabric-Wide Consistency Policy 11-12
fabricShow command 11-43, 2-24
fans, status of 14-9
fastwrite 21-16
FC routing 5-7
concepts 11-2
supported platforms 11-1
types 11-9
FC4-48 blade 9-1
FCAP 11-14
FC-FC routing 5-7
FC-FC routing service 11-1
FCIP 15-27
FCIP link 11-38
FCIP tunneling service 21-1
fcrConfigure command 11-13
fcrProxyConfig command 11-29
fcrProxyDevShow command 11-47
fcrResourceShow command 11-33
fcrXlateConfig command 11-29
fddCfg command 11-12
feature licenses 2-14
Fibre Channel
NAT 11-10
over IP 11-14, 21-1
routing 11-2
FICON 15-28
FICON environment
cascaded configuration 12-3
changing domain id 2-18
configuration settings 12-6
disabling IDID mode 12-4
displaying
link incidents 12-5
registered listeners for link incidents 12-10
enabling IDID mode 12-4
high integrity fabric 12-3
identifying port swapping nodes 12-11
monitoring FRU failures 12-10
node identification data, displaying 12-10
switched point-to-point configuration 12-3
switches, configuring 12-6
FICON-MIB 4-14
file copy, secure 4-23
filter-based monitors 17-8
adding 17-9
deleting 17-11
restoring configuration 17-16
saving configuration 17-16
firmware download 8-2
firmwareDownload command 15-30
fosConfig command 11-13
FR4-18i blade 9-1
frame transfer with brocade remote switch D-2
FreeRADIUS 3-22
front domain consolidation 11-5
FRU failures 12-10
monitoring in FICON environments 12-4, 12-10

G

gateway D-2
gateway, remote switch D-2
gathering information for technical support 15-3
GbE ports 9-2, 21-1
generating a public/private key 4-8
generating and storing a csr 4-9

H

HA failover 3-13, 3-22
HA-MIB 4-14
hard zoning 20-6
hardware-enforced zoning 20-6
haShow command 2-24
help information 1-3
high availability (HA) 2-24
high integrity fabric 12-3
HomeAD 3-20
host
    configuring 14-14
    reboots A-3
host-based zoning 20-3
HP/UX procedure A-15
HTTP 4-11
HTTPS 4-7, 4-11, 4-13
    certificates, security 4-2
hybrid update A-9

I

IAS (Internet Authentication Service)
    configuring users 3-24
    remote access policies 3-26
ID, account 2-3
identifying
    ports 9-1
    ports by port area ID 9-2
    ports by slot and port number 9-2
    ports from the tag field 12-20
    media-related issues 15-19
IDID 12-1
IDID mode, enabling and disabling in a FICON
    environment 12-9
IFL 11-2, 11-4, 11-10, 11-16
impact of changing fabric PID format A-2
implementing administrative domains 7-13
inaccurate information in the system message log 15-25
initializing trunking on ports 19-4
insistent domain ID 12-1
installing
    certificates 4-10
    root certificate to the Java plug-in 4-12
    switch certificates 4-10
interfabric link, *see also* ISL
    configuring 11-16
intermix mode 12-1
Internet Explorer 4-7
interoperability, switch requirements B-2
interopMode command 11-12
IP switch address 2-16
IP security 21-6
ipAddrSet command 2-6
IPComp command 21-6
IP-NAT 11-10
IPSec 21-6
    changeable parameters 21-9
    fixed parameters 21-8
ISL (interswitch link) 2-22, 21-1
    maximums 2-22

J

Java support, SSL 4-7
Java version 4-7

K

key transaction, for licensed features 2-14
keys, purchasing 2-19

L

legacy FCR switches 11-37
license ID 2-14
license key, activating 2-15
licenseAdd command 2-15
licensed features 2-14
licenseIdShow command 2-14
licenseRemove command 2-16

licenses
license ID 2-14
overview 2-14
purchasing keys 2-19
remove feature 2-15
transaction key 2-14
unlocking features 2-14
link incidents, displaying in a FICON environment 12-5, 12-10
link speeds in GbE port tunnels 21-3
linking through a gateway 2-23
Linux, configuring RADIUS 3-22
listing link characteristics 19-11
local accounts
disabling 3-9
enabling 3-9
local authentication, overview 3-30
LOCL (local clock) 2-14
logging timestamp 2-9
login
changing password 3-11
failure 2-2
switch 2-2
with Admin Domains 7-7
login sessions, maximum 3-2
long distance ISLs 18-2
LSAN 11-22
LSAN zoning 11-47
lsanZoneShow command 11-50
LUN sharing 11-38
LWL, ISL Trunking support 19-1

M

M_Port 15-27
maintaining
licensed features 2-14
zones 20-13
making basic connections 2-22
managing
accounts 3-11
zoning configurations in a fabric 20-23
mask for end-to-end monitors
displaying 17-7
setting 17-6
matching fabric parameters 11-11
McDATA 11-43
connecting 11-38
members
policy 5-2
policy, adding 5-9
policy, removing 5-9
merging zones 20-16
MIB 4-14, 4-16
mibCapability command 4-18
modifying zoning configurations 20-19
monitoring
end-to-end performance 17-3
filter-based performance 17-8
ISL performance 17-12
resources 11-33
traffic 19-5
trunks 17-12
monitors, clearing counters 17-15
Mozilla 4-7

N

names
chassis 2-17
switch defaults 2-16
NAT 11-10
network address translation, *see* NAT
network security 4-3
node identification data 12-10
nonfcsadmin role 3-21
normal operation, FC routing 11-8
NR_Port 11-34
nsAllShow command 2-25, 11-43, 11-49
nscCmShow command 11-49
nsShow command 2-25
NTP access 2-14
numbering ports 9-1

O

object connectivity event 4-22
obtaining and unzipping firmware 8-3
obtaining
  certificates 4-10
  slot information 9-13
offline update A-8
online update A-8

P

passwd command 3-11
passwdCfg command 3-13
passwords 2-3
  boot prom 3-30
  changing 3-11
  changing defaults 2-5
  default 2-4
  expiration policy 3-15
  limits 2-4
  management information C-1
  migration during firmware changes C-4
  policies 3-13
  PROM 3-30
  prompting behaviors C-3
  recovering forgotten 3-35
  recovery 3-33
  recovery options C-5
  recovery string 3-30
  rules 3-11
  strength policy 3-14
perfAddEEMonitor command 17-4
perfAddIPMonitor command 17-9
perfAddUserMonitor command 17-10
perfCfgRestore command 17-16
perfCfgSave command 17-16
perfDelEEMonitor command 17-8
perfDelFilterMonitor command 17-11
performance monitoring commands 17-2
performing PID format changes A-12
perfSetPortEEMask command 17-6
perfShowAlpacrc command 17-3
perfShowPortEEMask command 17-7
permissions and roles 3-3
persistently enabling/disabling ports 12-16
phantom domains 11-4, 11-10
physical fabric administrator 7-6
PID
  AIX procedure A-17
  binding A-1
  earliest supported firmware version A-10
PKI 4-7
planning update procedure A-7
platforms, FC routing supported 11-1
POD
  activating 2-20
  activating ports 2-19
  enabling ports 2-21
  full fabric access 2-19
policies
  account lockout 3-16
  activating 5-8
  adding members 5-9
  changes, activating 5-8
  changes, saving 5-8
  creating 5-4, 5-7
  creating, DCC 5-4
  creating, SCC 5-7
  DCC 5-4
  deleting 5-10
  members, adding 5-9
  members, identifying 5-2
  members, removing 5-9
  password expiration 3-15
  password strength 3-14
  routing 10-1
  SCC 5-7
port hardware state 4-22
port mirroring 15-27
port swapping nodes, identifying in FICON environments 12-11
port-based routing 10-2
portCfgExPort command 11-40
portLog command 11-47
ports 2-21
  activating 2-19
  activating POD 2-20
  buffer-limited 18-2
  enabling 2-21
  GbE 9-2
  identifying 9-1
  identifying by port area ID 9-2

identifying by slot and port number 9-2
licenses 2-19
numbering 9-1
status of 14-6
swapping 12-11
switch naming standards 12-17
portShow command 11-40
portStop command 11-40
powering off a card 9-7
powering port cards on and off 9-7
preparing a switch 12-7
private key 4-8
procedural differences between fixed-port and variable-port switches 1-1
PROM
password 3-30
recovery string 3-30
setting password 3-30 to 3-32
protocols
secure 4-1, 4-2
status 4-22
proxy
devices 11-8
PID 11-4, 11-29
public key 4-8
public key infrastructure encryption 4-7

R

RADIUS 3-21, 3-25, 3-27, 3-28, 3-29
RADIUS clients
switch configuration 3-23
Windows configuration 3-24
RADIUS server 3-17
configuration 3-22
configuring 3-22
LINUX configuration 3-22
RADIUS service
Windows configuration 3-24
RBAC (Role-Based Access Control) 3-2
reboot 11-34
recognizing
buffer underallocation 19-12
MQ-WRITE errors 15-13
port initialization and FCP auto discovery process 15-26
recording configuration information 12-22
recovering
accounts 3-10
forgotten passwords 3-35
recovery password 3-33
recovery string 3-30
boot PROM password 3-30
registered listeners 12-10
remote access policies 3-26
remote switch D-1
removing
Admin Domain members 7-17
alias members 20-12
end-to-end monitors 17-8
features 2-15
filter-based monitors 17-11
licensed feature 2-15
members from a zone configuration 20-20
zone members 20-14
renaming Admin Domains 7-18
resolving zone conflicts 20-29
restoring
configurations in a FICON environment 6-9
monitor configuration 17-16
segmented fabric 15-8
system configuration settings 6-5
Role-Based Action Control. *See* RBAC.
roles
domain defaults 2-4
for user accounts 3-2
nonfcsadmin 3-21
routing
assigning a static route 10-3
policies 10-1, 10-2
RSCN 2-16
rules
configuring zones, for 20-10
password 3-11

S

SAN Pilot 11-43, 11-48
saved zone configuration 20-6
saving
  monitor configurations 17-16
  policy 5-8
  policy changes 5-8
saving and restoring monitor configurations 17-16
scalability 11-38
SCC (Switch Connection Controls)
  list 11-15
  policy 5-7
scope and references 1-2
secCertUtil 4-8, 4-9, 4-10, 4-13
secure file copy, configuring 4-23
Secure Fabric OS 11-14
secure shell (ssh) 4-3
secure socket layer (SSL) protocol 4-7
secure telnet certificates 4-2
security 4-3
  activating certificates 4-11
  and tunneling 21-6
  Brocade MIB 4-14
  browsers 4-7
  certificates 4-2
  certificates, deleting 4-13
  certificates, displaying 4-13
  configuring standard features 4-1, 5-1
  enabling CHAP 3-25
  encryption 4-7
  FibreAlliance MIB 4-14
  file copy 4-23
  HTTPS, certificate 4-2
  IAS remote access policies 3-26
  obtaining certificates 4-10
  secure protocols, supported 4-1, 4-2
  secure telnet, certificate 4-2
  setting levels 4-15
  SNMP configuration 4-14
  SNMP default values 4-21
  SNMP traps 4-14
  SNMP values 4-20
  SSH, certificate 4-2
  SSL, certificate 4-2
  zoning 20-28
selecting a PID format A-4
sensor
  event 4-22
  status 4-22
sequential devices 21-16
serial connection 2-2, 2-3
server database, configuring 13-6
sessions, maximum allowed 3-2
setting
  boot PROM password 3-30
  card beacon mode 9-17
  chassis configurations 9-10
  default account passwords 2-9
  IP address 2-7
  mask for end-to-end monitors 17-6
  mode register bits 12-16
  port speeds 19-7
  security levels 4-15
  switch date and time 2-9
  time zone 2-11
  unique domain id 12-9
setting up
  automatic trace dump transfers 14-16
  RADIUS AAA service 3-21
settings
  changing passwords 2-5
  CHAP local security 3-25
  date and time 2-9
  PROM password 3-30 to 3-32
  security level 4-15
  SNMP 4-20
  SNMP default values 4-21
setup summary 12-12
SID 15-28
Silkworm 12000 and 24000 default accounts 2-4
SLAP 11-14
slotShow command 2-24, 11-11
SNMP 4-14, 4-15, 4-19, 4-20, 4-21
  ACL 4-1
  agent 4-1
  and password change 3-11
  attributes 4-15
  configuring 4-14
  MIB traps 4-22, 4-23
  support overview 1-2
snmpConfig command 4-14, 4-16, 4-17, 4-19
snmpMibCapSet command 4-14, 4-22
snmpMibCapShow command 4-14
SNMPv1 4-14, 4-17, 4-19

SNMPv3 4-14, 4-16
specifying
frame order delivery 10-3
routing policies 10-2
speeds, ISL 2-22
SSH
certificates 4-2
client, configuring 4-3
SSL 4-7, 4-8, 4-11, 4-13
certificates, security 4-2
SSL protocol, configuring 4-7
standard filter-based monitors 17-9
standard security features, configuring 4-1, 5-1
standard trunking criteria 19-2
standby CP blade 3-22
state
connection unit 4-22
port hardware 4-22
static PID mapping errors A-3
static route 10-3
status
connection unit 4-22
connectivity 4-22
protocol 4-22
sensor 4-22
storage-based zoning 20-3
summary
PID formats A-2
SSL procedures 4-7
support
FC router 5-7
Java version 4-7
SNMPv1, legacy 4-19
SNMPv3 and v1 4-14
supported
features B-2
protocols 4-1, 4-2
supportSave command 11-8, 14-15
supportShow command 11-8
swapping
port area IDs A-18
ports 12-11
SW-EXTTRAP 4-14
switch
access 4-5
access methods, CLI 1-2
access methods, Fabric Manager 1-2
access methods, Web Tools 1-2
certificates, installing 4-10
configuring 3-27, 12-6
configuring single 12-7
connecting 2-22
deleting RADIUS configuration 3-29
disabling 2-20
disabling ports 2-21
displaying RADIUS configuration 3-27
enabling 2-20
FICON environment, configuring 12-6
identifying 2-16
IP 2-16
name, customizing 2-16
name defaults 2-16
name limitations 2-16
RADIUS client 3-23
RADIUS configuration 3-27
RADIUS configuration, adding 3-28
RADIUS configuration, disabling 3-28
requirements for interoperability B-2
system status 14-4
user-defined accounts 3-7
Switch Connection Controls, *see* SCC
switchShow command 2-24, 11-42
SWL, ISL Trunking support for 19-1
synchronizing date and time 2-13
syslogd, configuring 14-13
system-defined Admin Domains 7-4
systemGroup command 4-18

T

tag field, interpreting 12-20
tape pipelining 21-16
tape write acceleration 21-16
TCP/IP 21-6
telnet
configuring 4-4
connection 2-2
support overview 1-2
temperature, status 14-10
time and date 2-9
time zones 2-9, 2-13
tools, CLI overview 1-2
tracking and controlling switch changes 2-25

traffic patterns
planning for 19-4
transaction key 2-14
traps
MIB 4-14
SNMP 4-14
SNMP MIB traps 4-22
troubleshooting 12-19
certificates 4-13
corrupt certificate 4-13
invalid certificate 4-13
migration 8-20
port mirroring 15-27
trunking problems 19-11
trunking
displaying information 19-8
over extended fabrics 19-9
tsClockServer command 2-11, 2-14
tsTimeZone command 2-10, 2-11
tunneling and security 21-6

U

unlocking licensed feature 2-14
unsupported features B-2
upgrading
directors 8-9
switches 8-7
user databases 3-13
user, local
disabling 3-9
enabling 3-9
userConfig command 3-3, 3-8, 3-9
user-defined
accounts 3-7
Admin Domains 7-4
filter-based monitors 17-10
users, authenticating 3-1
using
certificates 4-7
dynamic load sharing 10-4
FICON CUP 12-12
legacy commands for SNMPv1 4-19
legacy SNMPv1 4-19
zoning to administer security 20-28

V

validating Admin Domain members 7-20
VE_Port 15-27
vendor switch requirements B-1
verification check 11-11
verifying
device connectivity 2-22
fabric connectivity 2-24
high availability (HA) 2-24
version command 11-11
VEX_Port 21-3
viewing
alias 20-13
equipment status 14-9
fan status 14-9
port information 14-6
port log 14-11
port status 14-6
power supply status 14-10
power-on self test (POST)14-2
routing path information 10-5, 10-7
switch status 14-4
system message log 14-11
temperature status 14-10
zone database configurations 20-23
zones 20-14
viewing and saving diagnostic information 14-15
virtualized ports 21-2
VSA 3-17

W

Web Tools
access methods 1-2
support overview 1-2
Web Tools—AP Edition 11-26
Windows RADIUS, configuring 3-24
working with domain IDs 2-17
WWN 2-15, 2-16
Admin Domains 7-9

X

xlate domains 11-4

Z

zones
- adding members 20-14
- adding switches 20-26
- administering security 20-28
- aliases 20-5
- aliases, creating and managing 20-11
- concepts 20-1
- configuring rules 20-10
- creating 20-13
- database size 20-16
- default zone mode 20-15
- defined zone configuration 20-6
- deleting 20-14
- deleting configurations 20-21
- disabled zone configurations 20-6
- effective zone configuration 20-6
- enforcement 20-6
- hard zoning 20-6
- hardware-enforced zoning 20-6
- host-based 20-3
- LUN masking 20-3
- maintaining 20-13
- merging 20-16
- name restrictions B-5
- name server-based 20-3
- objects 20-5
- optimizing resources 20-1
- removing members 20-14
- removing members from a configuration 20-20
- resolving conflicts 20-29
- saved zone configurations 20-6
- schemes 20-5
- soft-zoning 20-3
- splitting a fabric 20-28
- storage-based 20-3
- terminology 20-1
- troubleshooting 20-29
- types 20-3
- viewing 20-14
- viewing configurations 20-23

zone configurations 20-6
- adding members 20-20
- creating and maintaining 20-19, 20-20

zoning
- LSAN 11-47
- PDCM considerations 12-18
- restrictions B-4